

ETHICAL AND EPISTEMIC EXPRESSIONS

Matthew Tate Chrisman

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Philosophy.

Chapel Hill
2006

Approved by:

Geoffrey Sayre-McCord

Jay F. Rosenberg

William G. Lycan

Jesse J. Prinz

Gerald G. Postema

© 2006
Matthew Tate Chrisman
ALL RIGHTS RESERVED

ABSTRACT

Matthew Tate Chrisman: Ethical and Epistemic Expressions
(Under the direction of Geoffrey Sayre-McCord)

My dissertation has two principal aims. First, I develop a novel account of the meaning of ethical claims that provides a way to retain the advantages while overcoming problems with prominent descriptivist and noncognitivist metaethical accounts. Second, I show how this account extends naturally into an account of the meaning of *epistemic* claims in a way that can recast and begin to adjudicate recent metaepistemological debate between prominent contextualists and invariantists. Thus, in my dissertation the two central paradigms of normative discourse receive parallel treatment.

Many philosophers treat ethical claims as descriptions of a special kind of fact, which raises difficult questions about the nature of ethical facts and our access to them. Other philosophers have argued that ethical claims are not descriptions but instead expressions of noncognitive attitudes, which avoids ontological and epistemological commitments but undercuts the apparent truth-evaluability and cognitive structure of ethical claims. My account of the meaning of ethical claims offers a third option, which I argue retains the benefits of descriptivism and noncognitivism while avoiding the drawbacks. To develop this view, I exploit the inferentialist idea that part of the meaning of ethical claims can be given in terms of what sort of commitment they express, where types of commitments are individuated by their inferential role. In brief, the thesis is that ethical claims express distinctively practical commitments.

Although there is much recent debate about the truth-conditions of knowledge claims, many epistemologists have not even considered the epistemological parallel to the question at issue between ethical descriptivists and noncognitivists: What is the expressive role of knowledge claims? In the second part of this dissertation, I develop an inferentialist answer to this question modeled on my inferentialist version of ethical expressivism. The idea is to treat knowledge claims as the expression of a particular kind of distinctively practical commitment—roughly, it's the commitment to trust someone's judgment about some matter.

PREFACE

When wandering through the corridors of Caldwell Hall on late-night study-breaks, one can be inspired to pick up dissertations written by people whose names are vaguely familiar from departmental lore. If you are like me, you ask yourself questions like: “What does a complete dissertation look like?” “Has Jay Rosenberg ever approved of any dissertation?” “What’s the shortest possible dissertation?” “She got a good job; what did she write her dissertation on?”

So you start reading. If you are anything like me, however, you rarely read past the abstract and preface. For this seems to carry some obligation to engage philosophically with the dissertation; the reflexes engendered by proto-seminar kick in and you begin to ask: what’s the main thesis, what’s the main argument, why is this important, where are the holes? But then you remember that you had picked up the dissertation in the course of *taking a break* from engaging philosophically with what was probably a more important and more clearly written piece of philosophy.

It is, therefore, unfortunate that so many dissertations either lack prefaces altogether or have prefaces that do barely more than thank the supervising committee. I think we should strive for more substantial prefaces. (Ironically, 2006 is the first year in which dissertations can be submitted electronically, so the likelihood of stumbling onto this preface while wandering through the corridors of Caldwell Hall will be quite low.)

As I rationally reconstruct it, this dissertation began its life in the summer of 1998 in a Mexican dive in Durham, North Carolina called the Cosmic Cantina. Somehow my friends Graham Hubbs and Jonah Johnson and I had convinced Duke University's Talent Identification Program to hire all of us as teaching assistants for various classes in their summer nerd-camp for high schoolers. We were often forced to seek respite at the Cosmic Cantina from the brutal (7 hours a day) teaching load and the AC-free and alcohol-free dorms.

One night, with a plastic picnic table full of burrito wrappers and Corona bottles in front of us, I insisted—despite appreciating the manifest punkness of their Nietzschean- and Foucaultian-inspired disdain for concepts like ‘moral objectivity’ and ‘moral truth’—that it is a desideratum of moral theory to construe many ordinary moral judgments as true. They balked: Moral discourse is nothing like ordinary discourse about objects of perception like tables and chairs—for it attempts to sway our wills. Moreover, moral discourse is nothing like scientific discourse—for morality lacks anything like the self-corrective methodology and explanatory power of science. Moral truths are a fiction—moral reality is a hoax. I'm not sure how boldly Graham and Jonah would assert this thesis today. But the vigor with which they defended it at the Cantina back then left a lasting impression on me.

I later learned that many metaethicists don't want to concede that moral discourse is interestingly different from paradigmatically descriptive discourses, but I was then (and today remain) convinced that there is such a difference. However, I was then (and today remain) convinced that it is a desideratum of moral theory to construe many ordinary moral judgments as true. Much of my work in metaethics has been driven by the goal of reconciling these two deeply held but apparently conflicting convictions. The key, as I see it, is to conceive of moral discourse as aiming at a distinctive kind of truth—practical truth. These are truths not about the way the world is but about the way to live in it. But, lest I scare off

the casual procrastinator even before the end of this preface, let me now merely promise much more discussion of that in the pages that follow. My present aim is to acknowledge the inspiration provided by Graham and Jonah. This dissertation is dedicated our friendship.

My parents have kindly supported and encouraged my intellectual development throughout my life. This dissertation may not make me the kind of doctor “that helps people”—at least not *physiologically*. However, I am grateful to my parents for helping me to realize that there are other ways to help people and that some of these involve asking the right questions rather than having the right answers.

I also want to express my appreciation to the people who supervised both this dissertation and, more generally, my graduate education. As I write this, the philosophy graduate program in Chapel Hill is thriving. This is due in large part to the extraordinary efforts of the person who has been chair of the department for my tenure: Geoffrey Sayre-McCord. I thank him for finding time in his busy schedule as departmental chair to be a good chair of this dissertation committee, to read my work carefully, and to offer insightful criticisms, suggestions, and amplifications. Geoff has been a constant presence in my graduate education—from his seminar on 20th Century Ethical Theory to his being on my M.A. committee to his meeting me weekly one summer to brainstorm ideas for my dissertation to his coaching me into the realization of some of these ideas. I owe a debt of gratitude to Geoff that one can owe only to one’s dissertation director. Hopefully what follows this preface is the first of many installments in my repayment.

Jay Rosenberg and William Lycan were the other two committee members on this dissertation, and both have been philosophical inspirations to me as well as incredibly diligent committee members. Jay successfully transmitted to me his admiration for the work of Wilfrid Sellars; moreover, Jay’s own writings have made Sellarsian themes and theses much more accessible to me. Early on in the project, he suggested that I take a look at

chapter 7 of Sellars' *Science and Metaphysics*, which is a chapter in metaethics. In the context of that book, the chapter almost comes across as an afterthought. But, as a result of Jay's suggestion, the positive metaethical view which is at the heart of this dissertation is inspired by the problematic Sellars sought to address in his brief foray into metaethics. It was in Bill's seminar on epistemic contextualism that I had and first articulated (in the form of a seminar paper) the leading structural idea of this dissertation. This was to see metaethical relativism as parallel to epistemic contextualism, and then to explore the expressivist alternatives to relativism/ contextualism with a clear eye on each side of the ethics-epistemology divide. Bill's encouragement of this idea has allowed it to flourish.

Despite being committed only as "readers" on this dissertation, Jesse Prinz and Jerry Postema have been very active in reading various drafts of the material and giving me helpful feedback.

As members of my M.A. committee, Mike Resnik and Keith Simmons were both seminal influences on my early work towards this dissertation. I want to thank Mike especially for his sage advice at several crucial stages in my graduate career.

Ever since he came to Chapel Hill, Ram Neta has been a seemingly indefatigable philosophical conversant with virtually all members of the philosophical community here. I want to thank Ram for suffering through many of my nascent paper drafts; many of the thoughts he forced me to have found their way into this dissertation. It was only after I was fairly deep into this project of defending an expressivist account of ethical and epistemic claims that I discovered Dorit Bar-On's fascinating and rigorous work on self-expression—see especially Bar-On (2004). I was pleased to learn how much our views on expression resonate; and I want to thank her for all that I have learned from her in the course of our collaborative work on a paper.

I have discussed the themes of this dissertation with many other people. I would like to express my appreciation especially to Stew Cohen, Jamie Dreier, Hartry Field, Tom Hill, David Landy, Nicoletta Orlandi, Ted Parent, David Reeve, Mike Ridge, and Piers Turner. A Dissertation Completion Grant from the Charlotte Newcombe Foundation insured that this dissertation was completed in a timely fashion; I thank the foundation for its support. An Off-Campus Dissertation fellowship from the UNC Graduate School and summer grant from the Deutscher Akademischer Austausch Dienst allowed me to spend two terms working on this dissertation in Tübingen, Germany; I thank these organizations for their support. In Tübingen, I was fortunate to have a connection to the Philosophisches Seminar in the grand old Alte Bursa, where frequent philosophical exchanges on a variety of topics with Manfred Frank, Anton Koch, and Frank Hoffman were a great pleasure. To them I express my gratitude for making my stay in Tübingen so philosophically engaging. My apartments on Damaschkeweg and Schleifmühlweg were ideal environments to pursue sustained writing.

I want to express my appreciation to Kristin Funcke for the reprieve from this writing that her companionship supplied. Without it I would have gone crazy, or at least not had a reason to learn to cook Eier mit Senfsoße.

Pursuing a dissertation in philosophy can seem to be a perilous endeavor. The job-market for philosophy Ph.D.'s is perhaps more open than it has been in the past 15-20 years, but it is still quite tight. One begins graduate school satisfied that, despite job-prospects, this is what one most enjoys doing, but then it becomes difficult, and long, and one doesn't make much money, etc. One is taught to be critically rigorous, but one is merely told to be "original" and "creative". One has to learn to love the austere landscape of a blank page, and then to fearlessly fill it with words, upon which others will exercise the "critical rigor" they have been taught. If things go right, though, there is a Gestalt-shift where one becomes both deeply invested in one's own views and convinced that the best way to make them better is

through the dialectical back-and-forth of characteristic of good philosophy. It was in the course of writing this dissertation that I experienced this Gestalt-shift. And so now it is with pleasure that I invite your critical rigor!

206A Caldwell Hall

Chapel Hill, North Carolina

13 February 2006

Motto to which I aspire: “Coming with clean concise thoughts, penetrating patterns...”

—Big Jus, from “The Fire in which You Burn”

CONTENTS

Introduction.....	1
-------------------	---

Part 1: Ethics

1. Descriptivist Accounts of Ethical Claims.....	14
I. Naturalism.....	16
II. Nonnaturalism.....	39
III. Relativism.....	55
IV. Three Challenges.....	70
2. Expressivist Accounts of Ethical Claims.....	73
I. Emotivism.....	75
II. Quasi-Realism.....	88
III. Inferentialist Expressivism.....	104

Part 2: Epistemology

3. Descriptivist Accounts of Epistemic Claims.....	149
I. Three Accounts.....	151

II. Skeptical Paradox.....	154
III. Ordinary Discourse.....	167
IV. Three Challenges.....	177
4. Expressivist Accounts of Epistemic Claims.....	180
I. Epistemic Noncognitivism.....	181
II. Epistemic Inferentialist Expressivism	191
 Bibliography.....	 221

INTRODUCTION

This dissertation has two principal aims. The first is to develop a new account of the meaning of ethical claims that is reactive to the drawbacks of previous attempted accounts. The second is to show how this metaethical account can be extended into a metaepistemological account of the meaning of epistemic claims. Thus, in my view *the* two central paradigms of normative discourse require and here receive recognizably parallel semantic accounts.

The metaethical account I defend here is at the same time expressivist and Kantian, which is a synthesis of what most philosophers have taken to be incompatible ways of thinking; but, in my opinion, precisely this synthesis is required to overcome problems besetting extant metaethical accounts. It is expressivist in that it stresses and turns on a difference in kind between what is expressed by ethical claims and descriptive claims.¹ Typically, expressivists have been Humeans of one sort or another, in that they see the key to understanding the meaning of ethical claims as understanding the kind of *psychological*

¹ To avoid confusion, it's perhaps worth mentioning that sometimes expressivism is understood as the thesis that ethical claims are "expressive" in contrast to empirical claims which are descriptive. I think this is a confused way of understanding the expressivist's characteristic appeal to the fact that on their view ethical claims express some mental state but do not describe one as being in it (as for example an ethical subjectivist would). In my view, *all* claims express *something*, and the issue is whether ethical claims express an interestingly different kind of thing (usually, mental state) than descriptive claims.

state they express; specifically, expressivists have tended to defend some version of the thesis that unlike descriptive claims which express beliefs about the way the world is, ethical claims express nonbelief attitudes with a special motivational character. In contrast, my account is Kantian in that I think that the key difference between ethical and descriptive claims is best explained not in terms of an empirical distinction between two types of psychological states, but in terms of an antecedent normative distinction between two types of reasoning—theoretical and practical.²

One surprising but extremely satisfying upshot of the development of this new version of metaethical expressivism is that it opens up room for an expressivist alternative in and, in some ways, the very idea of a “metaepistemology”. Recently much epistemological theorizing has been given over to analyzing the meaning of knowledge attributions with an eye to dispelling skeptical paradoxes and generally explaining the linguistic phenomena surrounding the verb ‘knows’. This has been called the “new linguistic turn in epistemology”, and to the reflective observer versed in the history of 20th century metaethics, the parallels to the old linguistic turn in ethics are remarkable—we see philosophers temporarily bracketing base-level normative-epistemological issues and defending the meta-level views that (i) knowledge is simple and unanalyzable, (ii) knowledge attributions are speaker-relative, (iii) knowledge attributions are subject-relative, (iv) knowledge attributions are uniformly in error. Most remarkable of all is that no one has yet tried to run the expressivist gambit of arguing that the problem with all of these views is that they tacitly assume that knowledge attributions are straightforwardly descriptivist claims and one can develop a better account

² Of course, it’s conceivable that the empirical distinction between types of psychological states can be explained in terms of the antecedent normative distinction between types of reasoning. This is not the traditional expressivist order of explanation, since expressivists typically want to reduce normative distinctions to psychological distinctions in order to “naturalize” them; however, arguably this normative understanding of the psychological has been the thrust of some very recent work by Blackburn (1998) and Gibbard (2003). In my view, we should go straight to the normative distinction in our metaethical theorizing.

of knowledge attributions by exploring the way they're expressively different from descriptivist claims.³ I suspect this oversight is due to the fact that it seems patently absurd to suggest that knowledge attributions don't express beliefs. With my Kantian form of metaethical expressivism on the table, there is room go a different direction by arguing for the thesis that knowledge attributions express practical rather than theoretical commitments (which both might exhibit an explicable connection to the psychological category "belief"), and thus generate a metaepistemological expressivist position with the resources to overcome many of the problems besetting other extant attempts to analyze knowledge attributions. Or so I shall here try to demonstrate by example.

First some stage setting is in order. My argumentative strategy in pursuit of the first principal aim of this dissertation is roughly dialectical. I think that ethical expressivism emerged on the metaethical scene in the 1930s as the result of two related pressures. The first was dissatisfaction with the moral ontology and epistemology implied by the sort of ethical nonnaturalism made popular at the time by G. E. Moore's open-question argument and insistence on the naturalistic fallacy. This pressure encouraged some metaethicists to argue that Moore's open-question argument does not establish that ethical facts are nonnatural facts because it ignores the possibility of *a posteriori* reduction of ethical facts to natural facts. However, the second pressure was dissatisfaction with the popularity of ethical relativism resulting from the boom of new anthropological and ethnological data attesting to the diversity of different ways of living. And this pressure combined with the first pressure

³ This is perhaps too quick. Sellars (1956) notoriously articulates and defends the thesis that knowledge attributions are claims made in the "logical space of reasons" and draws an analogy between the problems with attempts to reductively analyze knowledge to an epistemic version of the naturalistic fallacy articulated by Moore (1903). However, it is debatable whether Sellars should be construed as a proto epistemic expressivist; one might want to construe him as a proto epistemic nonnaturalist, à la Moore. In any case, insofar as he is a proto epistemic expressivist, my metaepistemological views (and, for that matter, my metaethical views) are certainly deeply indebted to Sellars.

encouraged other metaethicists to argue that Moore's open-question argument, the relativist's position, and the *a posteriori* reductive naturalist's position all treat ethical claims as descriptive of some realm of fact, but if we construe ethical claims as nondescriptive and thus nonreferential we can eschew the ontology and epistemology of the nonnaturalist while at the same time rejecting the views that cross-cultural ethical discussions and evaluations are spurious and that ethical facts reduce to natural facts.

In chapter 1, I review and distill this dialectic by arguing that each of the three most prominent descriptivist accounts of the meaning of ethical claims—a posteriori reductive naturalism, nonnaturalism, and relativism—are reactive to important considerations but at the same time suffer very serious objections. This generates three initial challenges to the development of an account of the meaning of ethical claims, which I call *Mackie's challenge*, the *action-guiding challenge*, and the *relativist's challenge*. I then take these into chapter 2 where I discuss traditional and contemporary expressivist alternatives to the descriptivist accounts. There I argue that although these various versions of noncognitivism can more or less meet the three challenges, even the best expressivist accounts suffer a very serious problem emerging from attempts to combine expressivism and minimalism. I use this to generate a fourth challenge, which I call the *challenge of creeping minimalism and valid reasoning*. It is particular to the development of a nondescriptivist account of the meaning of ethical claims. So I develop an alternative, inferentialist, form of ethical expressivism that I think can meet all four challenges.

The core idea behind my alternative is a rejection of the psychologistic way of explaining the meaning of ethical claims implicit in noncognitivist forms of expressivism in favor of an inferentialist way of explaining this meaning. Typically expressivists argue for a distinction between different kinds of psychological states then they argue that what speakers mean in making claims (utterance-meaning) is to be explained in terms of what psychological state is

expressed by these claims; and finally they suggest that the meaning of the sentences with which these claims are made (sentence-meaning) is to be explained in terms of the speaker's meaning in making them and thus in turn in terms of the psychological state expressed. Although I think there is a difference between what is expressed by ethical and descriptive claims that is key to understanding the distinctive meaning of each, I think this typical psychologist *cum* expressivist way of proceeding is wrongheaded for two related reasons. First, I doubt that sentence-meaning can really be explained in terms of what psychological states speakers mean to express by their multifarious uses of the relevant sentences. Second, the only plausible way such an explanation gets off the ground, it seems to me, is if the operative psychological terms themselves already carry some semantic import—are already recognizable intentional—however, if the expressivist assumes that they do, this only pushes the question back to explaining that import. These are not particularly novel criticisms of the psychologist accounts of semantic content, but I don't know of anyone who has brought them to bear on the psychologist order of explanation implicit in traditional and contemporary forms of ethical expressivism.

If one wants to retain the expressivist idea that a distinctive and illuminating feature of ethical claims is that they express something of an interestingly different kind from descriptive claims, something “nondescriptive”—which I do—these problems with the psychologist order of explanation encourage pursuit of an alternative order of explanation. I work from Sellars' suggestion that we can only understand what speakers express with their acts of claiming in relation to the possibility of using sentences to make these claims, which in turn can only be understood in relation to the possibility of giving and demanding reasons for these claims. Thus, for me unlike most card-carrying expressivists before me, ‘what is expressed by ethical claims’ is not treated in empirical psychological terms but in normative inferentialist terms. The animating idea is Kantian: a precondition making any genuine ethical claims at all is that, in addition to operating within a theoretical conception of the

world (including us and our place in it), we must also operate within a fundamentally different practical conception of our relation to the world. Specifically, I argue that there are two different kinds of reasoning corresponding to two different kinds of reason one can give for one's claims; and I use this distinction to mark out two different kinds of things that can be expressed by claims in general—viz. theoretical and practical commitments.⁴ Thus, at the end of my metaethical story here, the meaning of ethical claims is distinguished from the meaning of descriptive claims in that they express a special sort of practical commitment rather than a theoretical commitment. And this, I argue, is an account that meets the three challenges which emerge in chapter 1 in a way that also meets the fourth challenge facing traditional and contemporary forms of expressivism which emerges in the middle of chapter 2.

So, the plan for part 1 is to articulate and motivate a version of metaethical expressivism that is novel in that it explains the difference between what is expressed by ethical and descriptive claims in inferentialist rather than psychologistic terms. In part 2, my argumentative strategy in pursuit of the second principal aim of this dissertation is both dialectical and analogical. Here it will be helpful to pause to sketch four different semantic explanations one can provide for apparent linguistic disagreement, in order then to bring the animating analogy into view.

At a high enough level of abstraction, both the metaethical and the metaepistemological views developed in this dissertation stress a particular sort of linguistic phenomenon that I believe is characteristic of specifically normative discourse. This phenomenon is a theoretically interesting species of the quite general and mundane phenomenon that for

⁴ Although Hare (1952) doesn't put his leading thesis in terms of the commitments expressed by claims, he was perhaps the first to use this general strategy of using the distinction between two different kinds of reasoning to generate an expressive distinction between different types of claims.

many sentences ‘S’, we can find people who affirm ‘S’ and people who deny ‘S’—sometimes even a single person will affirm ‘S’ on one occasion and deny ‘S’ on another occasion. Call this the *phenomenon of apparent linguistic disagreement*; initially it invites three types of descriptivist explanations. First, sometimes a straightforwardly *universalist* account is correct. That is, affirming ‘S’ is understood to commit its speaker to the fact that-S, and denying ‘S’ is understood to commit its speaker to the fact that not-S, and since these commitments are inconsistent, one of the two speakers is mistaken. A universalist response is most plausible when there is good reason to think that there is a fact of the matter about whether S, and the mistaken party to the dispute lacks some piece of crucial evidence; for example: “The election was rigged.” Second, sometimes a *speaker-relative* account is correct. That is, the sentence ‘S’ is understood to have a deictic element in its logical form, such that in the context in which ‘S’ is affirmed it is true but in the context in which it is denied it is false, which means that neither speaker is mistaken. This is most plausible when there is an explicit indexical in the sentence: for example: “Milford is not here.” Third, sometimes a more subtle *subject-relative* account is correct. That is, affirming ‘S’ is understood to commit its speaker to the truth of some proposition p, which itself references some situation, and denying ‘S’ is understood to commit its speaker to the falsity of a similar but slightly different proposition p', which itself references some different situation, and since it possible for both p and not-p' to be true, neither the affirmation nor the denial of ‘S’ need be thought of as mistaken. This is plausible when there is some reason to think that each of the sentences is implicitly sensitive something about the practical or conventional situation of some particular person; for example “Gottfried’s smoking hash is illegal.”⁵

⁵ These three views are the most intuitive descriptivist responses to the phenomena; however, they do not exhaust all of the options that have been suggested. MacFarlane’s (2005) and Richard’s (2004) semantic relativism treats the semantic content of a claim as relative to the context from which it is *evaluated*, so the apparent disagreement could, on their views, be explained as a result of conflating two different contexts of evaluation. However, I don’t find these proposals generally plausible or plausible for the cases that will constitute my analogy, so I won’t discuss them further here.

The mundane phenomenon of apparent linguistic disagreement has a more interesting subspecies that I believe is characteristic of specifically normative discourse. This is that for a class of sentences ‘S’, we can find people who affirm ‘S’ and people who deny ‘S’, such that they seem to be disagreeing but there is no uncontroversial evidence that would resolve the dispute. By “uncontroversial evidence”, I mean facts that both sides could accept as facts, which render either S or not-S more likely to be true. This phenomenon is made familiar in popular culture by the abortion debate. One side says, “Abortion is wrong” and the other side says, “Abortion is not wrong,” and appeal to putative facts such as that life begins at conception or that women have an inalienable right to determine what happens to their bodies are, naturally, unlikely to resolve the issue because they will only “count” as facts for one or the other side of the debate.

Can any of the three truth-conditional accounts of apparent linguistic disagreement explain this subspecies of the phenomenon of apparent linguistic disagreement? In the case of ethical claims, each approach certainly has its defenders; however, there are serious *prima facie* problems with each approach. Universalists treat either the affirmation or denial of ‘S’ as mistaken, which makes it incumbent on them to explain what the mistaken party is missing. Since, however, there is no uncontroversial evidence to appeal to, we can imagine universalists on both sides trying valiantly to diagnose the mistake of the other side but failing. This suggests that there may be no fact of the matter about who is right, which, in different ways, is the position of the speaker-relativist and the subject-relativist. A speaker-relativist will treat ‘S’ as covertly indexical such that the affirmation can be true when indexed one way and the denial can be true when indexed another way. A subject-relativist will treat the affirmation of ‘S’ as asserting some proposition p’ and the denial of ‘S’ as asserting the negation of some different (although similar) proposition p”, and so not really in disagreement. The difficulty with both of these accounts of apparent disagreements comes in explaining, to put it bluntly, why the affirmations and denials of ‘S’ *appear to disagree*. If

speakers and audiences take the affirmations and denials to stand in conflict, then one will need to explain this semantic mistake.

It is of course possible that these *prima facie* difficulties can be overcome for one or another of the descriptivist accounts; however, expressivism can be seen as a nondescriptivist fourth alternative that deserves our attention. Universalism, subject-relativism, and speaker-relativism all seem to privilege the content expressed by the relevant *sentences* over the issue of what the *speaker* who utters the sentence is expressing. Expressivists, on the other hand, insist that the content of all sentences is to be explained in terms of, or at least in conjunction with, what speakers express in uttering the sentences. Here it is crucial to notice that the expressivist is operating with two different expression relations. First, there is a semantic relation between the sentence uttered and its semantic content—in the case of indicative sentences, a proposition. Second, there is what is by the expressivist's lights a more fundamental relation between a person's act of uttering a sentence and whatever it conveys to the audience.⁶

When it comes to sentences that are typically used to express a psychological state whose direction of fit with the world is descriptive, such as a descriptive belief, it is easy to blur the distinction between these expression relations. This is because it is natural to think of descriptive belief as an attitude towards a semantic content that is usually expressed by a sentence whose utterance would express the belief. The situation is importantly different when it comes to utterances that are typically used to express a state whose direction of fit with the world is nondescriptive such as an intention or desire. Here, although one can express these states by uttering an indicative sentence, they do not manifest the same

⁶ The phenomenon of implicature displays perhaps another expression relation. For example, the sentence, "He's an above-average student" usually conveys that he is not a superb student. It is wrong to say that the *sentence* expresses this proposition, but it might be right to say that an *utterance* of the sentence expresses this proposition. The expressive phenomenon I focus on here is not implicature. For although implicature is interesting, I don't think it helps to account for normative discourse *per se*. (For more on the nature of this expression relation, see chapter 2.)

relationship between the nature of the psychological attitude and the semantic content of the sentence whose utterance would express it. (Consider the sentence “I’ll go to the movies”—this sentence expresses a proposition about the future, but its utterance usually expresses an intention rather than a predictive belief.)

By mobilizing this complexity, expressivists have the resources to offer a fourth account of apparent linguistic disagreement. For some cases of apparent linguistic disagreement, they can mirror one of the other three accounts by arguing that the relevant claims express the speakers’ beliefs, whose content can then be given in universalist, subject-relativist, or speaker-relativist fashion. However, for the subphenomenon I’ve claimed to be indicative of normative disagreement, expressivists will typically argue that the relevant claims express a nonbelief attitude that nonetheless exhibit a sort of practical disagreement. Here, because these attitudes exhibit a different direction of fit with the world, there need be no fact or piece of information that either speaker is missing in order to see them as nonetheless disagreeing. They are, to use Gibbard’s (2003) phrase, “disagreeing about what to do.”

Now, in one sense the expressive phenomenon exploited by expressivism in its account of the apparent disagreement with normative sentences is familiar and doesn’t obviously involve a novel account of the meaning of ethical claims. We often use descriptive sentences to express nondescriptive psychological states. For example, at dinner, one might say, “I’d like more potatoes,” to express a desire for more potatoes. However, this pragmatic issue about the state of mind expressed by these sentences doesn’t clearly entail anything nonstandard about the meaning of the sentences. That is, one could recognize the differences between the expressive force of these utterances and utterances of the same sentence to express descriptive beliefs, and yet maintain a descriptivist account of the meaning of these sentences. Typically, however, expressivists see the theoretical situation in an importantly different way with respect to ethical claims. They argue that the distinction

between the expressive force of ethical claims and descriptive claims should be transmuted into our account of the semantic content of the sentences.⁷ And this is obviously in line with the general psychologistic order of explanation implicit in extant forms of expressivism.

However, even if we replace the psychologistic order of explanation of semantic content with the inferentialist order of explanation—as I propose in chapter 2—there remains an important sense in which what is expressed by particular sentences cannot be explained in abstraction from what is expressed by speakers who utter them. So, in either case, the pragmatic difference in what is expressed by utterances of a sentence makes a difference in the kind of semantic content the sentence can be thought to have. And this added level of complexity is the key to the alternative explanation of apparent linguistic disagreement available to expressivists.

Now it is possible to bring the analogy crucial to my argumentative strategy in chapters 3-4 into view: Recent epistemological theorizing has been obsessed with what we might call the semantic approach to skeptical paradoxes. The idea is that what is troubling about skeptical arguments is that they more or less convince us of a conclusion of the form “S does not know that p,” whereas when we forget about the skeptical argument and go on with our normal lives we will freely assert, “S knows that p.” This is obviously an instance of the phenomenon of apparent linguistic disagreement. Moreover, when we think about it, there seem to be many ordinary occasions (i.e. occasions not involving skeptical arguments) where one person in one situation will attribute knowledge and another person in another situation

⁷ Compare Gibbard: “In a different and more traditional sense, to be sure, one could treat the belief that I am about to pack and a decision to pack as having the same “content,” *that I will forthwith pack*, toward which I take different propositional attitudes: belief, and deciding to actualize. The two attitudes have different “directions of fit”, we can say, toward the same item of “content”. Here, though, I am opting to transmute force into content: I speak of a single attitude “accepting” that one can take towards distinct items of content, *that I will forthwith pack* and to the plan *to pack*”(2003: 47).

will deny knowledge. The traditional view about this instance of apparent linguistic disagreement implicit in much epistemological theorizing is in essence a form of universalism—although in the epistemological debate it goes under the label *invariantism* (Unger 1984). Recently this traditional view has come into question and several philosophers have forcefully argued for a version of speaker-relativism vis-à-vis knowledge attributions—although in the epistemological debate it goes under the label *contextualism* (Cohen 1988, DeRose 1992, Lewis 1995). And even more recently than that, several other philosophers have contended that a form of subject-relativism better explains the linguistic phenomena surround knowledge attributions—although in the epistemological debate it goes under the label *subject-sensitive invariantism* (Hawthorne 2004, Stanley 2005).

So in recent epistemological theorizing we see each of the descriptivist responses initially suggested by the phenomenon of apparent linguistic disagreement. Delineating the advantages and disadvantages of each of these accounts is the project I pursue in chapter 3. This generates three initial challenges to the development of an account of the meaning of knowledge claims, which I call *Austin's challenge*, the *challenge of the paradox*, and the *challenge of unexplained semantic blindness*. In my view, none of the prominent descriptivist accounts fully meets all three challenges.

What is important to notice here is that the current state of metaepistemological debates is exactly parallel to the state of metaethical debates when we restrict it to the three most prominent descriptivist accounts of apparent linguistic disagreement. This—in tandem with the idea that epistemological discourse is normative—naturally encourages exploration of an expressivist account of knowledge attributions. However, I think this line of investigation is immediately squelched if one thinks that the only way to be an epistemic expressivist is to be an epistemic noncognitivist—for knowledge attributions seem to be among the *paradigmatically* cognitive claims. Yet, the inferentialist expressivist account of ethical

claims developed in the second half of chapter 2 provides new hope for epistemic expressivism; and, therefore, in chapter 4, I articulate and defend an inferentialist expressivist account of knowledge attributions that is parallel to the expressivist account of ethical claims pursued in chapter 2.

The core idea is to treat knowledge attributions as expressive not of theoretical commitments about whether some subject stands in the knowledge relation to a particular fact but as expressive of a special sort of practical commitment. Again this difference is articulated not in the psychological terms characteristic of noncognitivist expressivist accounts but in inferentialist terms, and this makes my version of epistemic expressivism much more plausible than the extension of any existent form of ethical expressivism to the epistemic realm would be. (Among other things, this is what helps my account to overcome the epistemic analogue of the challenge of creeping minimalism and valid reasoning and what lets it comport with the pervasive intuition that epistemic claims are fully cognitive.) In a slogan, then, the plan for part 2 is to develop a version of epistemic expressivism that better explains the linguistic phenomena addressed by existent metaepistemological accounts while avoiding obvious problems for a noncognitivist expressivist account of knowledge attributions. And the upshot of parts 1 and 2 together is the framework for a satisfyingly unified inferentialist expressivist account of ethical and epistemic claims.⁸

⁸ Although I will for the most part restrict myself to ethical value claims such as “Magnanimity is good” and knowledge attributions such as “Peter knows that Suzy is home” and not delve into the complicated issue of how other sorts of ethical and epistemological claims are related to these.

1 DESCRIPTIVIST ACCOUNTS OF ETHICAL CLAIMS

Many philosophers take a form of descriptivism to be the default framework for understanding the meaning of any claim. That is to say that they think that the meaning of any claim is to be understood in terms of what fact it describes, or, as it is sometimes put, what fact it states. So a natural way to begin our search for an adequate account of the meaning of ethical claims is by construing them as aspiring to state ethical facts. Although we might differentiate between different possible accounts based on the handful of different ways general descriptivist semantics have been articulated and defended in recent decades, I think the debate here is most tractable if we abstract away from these differences and focus on descriptivism in general.

Descriptivism about ethical claims is typically pursued by moral realists of one stripe or other. Indeed, it can seem that the only way to be a descriptivist anti-realist is to hold an error-theory about ethical discourse: the thesis is that ethical claims aim at stating ethical facts but uniformly fail. Notoriously, this is the position defended by Mackie (1977: ch. 1; see also 1946), whose so-called *argument from queerness* can be summarized in the follow five points: (1) Ethical claims aim to be “objectively true”—that is, their truth does not depend on any desire or preference or policy or choice of particular groups of people. (2) Ethical claims are “intrinsically action-guiding” in that they commit one to the existence of a reason for action whose recognition can be intrinsically motivating. From (1) and (2), it (putatively) follows that (3) one

who makes an ethical claim is thereby committed to the view that there exists a corresponding fact “that involves the call for action or for the refraining from action, and one that is absolute, not contingent upon any desire or preference or policy or choice, his own or anyone else’s.” However, (4) If there were such facts, then “they would be entities...of a very strange sort, utterly different from anything else in the universe. Correspondingly, if we were aware of them, it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing anything else,” which is sufficient to undermine the view that there exist such facts. Thus, he concludes (5) all positive ethical claims are false.

Most think Mackie’s conclusion (5) is nonnegotiable false. We *all* know that that cutting off babies’ eyelids for fun is morally impermissible. Knowledge is factive, so it must be true that cutting off babies’ eyelids for fun is morally impermissible. So the challenge Mackie’s argument poses for any account of the meaning of ethical claims is to show where it goes wrong. It will prove useful to consider the premises out of order. So, in the following chapter, I’ll explore the possibility of denying premise (3)—this is a consideration of expressivism. In this chapter, I shall begin in section 1 with recent versions of metaethical naturalism. This is roughly the view that ethical facts are natural in the sense that are identical to or appropriately supervene on facts that admit uncontroversially of empirical investigation, which implies that ethical facts have no peculiar intrinsic action-guiding character but can influence what is to be done, like any other natural fact, via the appropriate uptake into practical reasoning. This involves the denial of (2). Then, in section 2, I shall consider recent versions metaethical nonnaturalism. This is roughly the view that ethical facts, although metaphysically and epistemologically *sui generis*, are not too queer to be admitted in our conception of reality. This involves the denial of (4). Finally, in section 3, I shall consider recent versions of metaethical relativism. Sometimes metaethical relativism is seen as an anti-realist or error-theoretic view because it denies the existence of universal ethical facts; however, I think it is best seen as a descriptivist view that construes ethical claims as sometimes successfully describing relative facts. This involves the denial of (1).

I think all of the descriptivist views are ultimately unsatisfactory. My aim in discussing them is two-fold: First, I think each is motivated by important although vague considerations that every view about the meaning of ethical claims needs to take into account. So I want to identify and try to reach some precision about these considerations. Second, I think each of the views suffers drawbacks serious enough to provide a presumptive case against it and a good litmus test for any account of the meaning of ethical claims. So I want to also identify and try to reach some precision about these drawbacks. In the end, this discussion should provide the resources to outline three complex challenges to developing an adequate account of the meaning of ethical claims, which I plan to take with me into the following chapters.

I. NATURALISM

In the heyday of logical-positivism, ontological commitment to all sorts of facts began to seem problematic because it could not be shown how the existence of these facts could be verified. This, in turn, made it questionable how we could have any reason at all to believe that the facts are part of the natural world; and an appropriately enlightened conception of reality seemed to require that one recognize the existence of only the natural world. Of the many facts that came to seem problematic, perhaps the most troubling were ethical facts. Ghosts and witches had always seemed like the product of superstition. Aesthetic facts are might be recast in subjectivist lights without violating ordinary intuitions too much. The gods had been implicitly dying in the imagination of the enlightenment long before Nietzsche noted that the most important one was dead enough to start stinking. However, when it came to ethical facts about how it is morally right and wrong to interact with one another, intuitions about these seemed incorrigible enough to turn almost any argument against them on its head—especially an argument based on something so obscure as the verificationist theory of meaning!

There is, however, a compatibilist way out of this dilemma: reinterpret the menu of normative ethical theories—e.g. utilitarianism, egoism, deontology, etc.—as attempts to analyze the meaning of ethical claims in naturalistic terms. Assuming that one of them is right, ethical claims semantically reduce to claims about natural facts, and the legitimacy of ethical discourse can be squared with a commitment to naturalism. Moore (1903) identified the biggest problem with this compatibilist way out of the dilemma: all naturalistic analyses of the meaning of ethical claims seem, upon reflection, patently false; for we can always ask without apparent conceptual confusion whether something having the natural property referred to by the *analysans* really has the moral property referred to by the *analysandum*. But if a claim that something has a particular ethical property is supposed to just mean the same as the claim that it has some natural property, such questions should be conceptually closed.

I'll discuss the positive account Moore drew from this argument in more detail below, but here it's important to note that contemporary metaethical naturalists such as Sturgeon (1985), Lycan (1988), Boyd (1988), and (Brink 1989, 2001) all accept Moore's rejection of synonymy-analyses of ethical claims in nonethical terms—this is part of a broader rejection of a priori conceptual analysis as a viable philosophic methodology. However, against Moore they object that that even if all synonymy-analyses of ethical *claims* fail, it is highly implausible to suppose that this is probative with respect to the nature of the *facts* putatively described by ethical claims. Specifically, they argue that the issue of the naturalness of ethical facts can be settled a posteriori—after all, the nature of water, heat, and even perhaps pain are all things that can be settled a posteriori.

What a posteriori naturalism has going for it is that it provides a template for treating ethical facts as perfectly respectable and admissible into a naturalistic ontology. But even if it would be nice to be able to treat ethical facts as a species of natural facts, why should we think that we can? Initial debate in this area took it that in order to legitimately treat ethical facts as universal

natural facts, we must show how they can be *identified* a posteriori with prosaically natural facts. And optimism in this project stemmed from two rather different sources.⁹ First, in the philosophy of mind, materialism—i.e. the view that mental properties are type-type identical to material, or more specifically, bio-physical properties—received new support from Smart (1959). He argued that, despite the fact that there is a difference between the concepts ‘sensation’ and any ‘brain state’ that might be claimed to be identical to sensation, this doesn’t imply that sensations cannot be type-type identifiable with brain states. For there are all sorts of type-type reduction in the natural sciences that do not turn on or require semantic synonymy. For instance, no one will argue that the terms ‘temperature’ and ‘mean molecular kinetic energy’ are synonymous, but arguably physical theory has managed to type-type reduce temperature to mean molecular kinetic energy. Likewise, one might suggest that a similar type-type reduction of the mental to the bio-physical is on the horizon in the cognitive sciences.

The second source of optimism comes from developments in the philosophy of language, where Kripke (1980) and Putnam (1975) developed a conceptual way to argue for a posteriori identities. This consists of two core moves. First, they argue that through “twin-earth” type thought experiments, we can establish whether a term is a rigid designator. The idea is to consider one of our actual terms ‘t’ and then ask ourselves whether the orthographically, phonetically, and pragmatically identical term as used on an imagined possible world twin-earth (or even just in imagined different country in the actual world twin-USA) would mean the same thing as ‘t’, if everything is basically the same on twin-earth (or in twin-USA) except that the nature of the referent of ‘t’. Many philosophers find it highly plausible that names such as ‘Nixon’ and natural-kind terms such as ‘water’ are rigid designators. This is because they think that ‘Nixon’ and ‘water’ have different meanings as used on earth and twin-earth (or in USA and twin-USA). This view is based on putatively ordinary counterfactual intuitions such as “If the

⁹ In the following two paragraphs, I have benefited from Timmons’ (1999: 40-46) discussion of these issues.

person who was the 37th president had actually been an alien calling himself Nixon, it wouldn't have been Nixon," or "If the stuff in our lakes and streams hadn't been composed of H₂O, then it wouldn't have been water." These intuitions indicate that earthlings and twin-earthlings (or Americans and twin-Americans) aren't talking about the same thing. Second, Kripke and Putnam argue that from this it follows that scientific investigation of the nature of the referent of a rigid designator provides a posteriori evidence for an identity claim. For example, chemical analysis tells us that the substance we refer to with the term 'water' is H₂O, and since 'water' is a rigid designator, if some substance is not H₂O, then it is not water; thus, water=H₂O and all facts about water can be reduced to facts about H₂O. In this way, these philosophers think that the meaning of our term 'water' is determined by what stuff actually causally regulates its use and thereby endorse a "causal theory of reference".

For my discussion here, the key point about the argument for materialism in the philosophy of mind and the argument for a posteriori identities in the philosophy of language is that they both turn on a fact that should have been obvious to even Moore—viz. that predicate synonymy isn't required for property identity. This opens up conceptual space for a new version of metaethical naturalism that doesn't turn on the viability of synonymy-analyses of moral claims in naturalistic terms and where one can treat ethical predicates as picking out a property that is identical to that picked out by some nonsynonymous nonethical predicate. Such an "a posteriori naturalism" avoids Moore's open-question argument, but now the pertinent question is whether there is not only conceptual space for the position but also good reason to think the position is correct. That is to say that even if it is to be achieved a posteriori, a posteriori naturalists need a story to justify the identification of ethical properties with natural properties. Here I think naturalists face serious difficulties.

Lycan pursues something like the Kripke-Putnam strategy: "I am inclined to think that moral expressions are rigid rather than flaccid designators...wrongness, goodness, and so on do

not seem to be 'role' concepts whose occupants are the objects of scientific research. But if moral predicates are rigid, then what is a plausible epistemology for the identification of their referents with natural properties?"(1988: 205-206). His answer is that we have modal intuitions about moral properties much like the modal intuitions about natural kinds exploited by Kripke's and Putnam's arguments. They are not the result of purely linguistic competence; however, as the testing of normative ethical theories through consideration of counterfactual thought experiments exemplifies, we do have the intuitions. And he argues that in both cases these modal intuitions support supervenience claims to the effect that if something didn't have some property P_1 , then it couldn't have the property P_2 . In both cases, these modal intuitions are fallible, but, according to Lycan, they emerge as our total theory of reality is developed and many of them seem to survive the process of theory building through wide-reflective equilibrium. And he follows Armstrong (1982) in suggesting, "that in most cases the best explanation of a supervenience relation is a type-identity hypothesis...Thus in particular, if moral properties are found by modal intuition to supervene on facts about utility, harm, degradation, or the like, we would be justified in identifying them with complex natural properties involving those things"(Ibid.: 206-207).

I think it is doubtful that ethical terms are rigid designators in a way that can underwrite a posteriori identity claims. Timmons (1999) argues that they cannot be because when we run the twin-earth thought experiments analogous to the ones used by Kripke and Putnam to establish that names and natural kinds terms are rigid designators, they seem to come out the wrong way. His argument proceeds as follows. Because there is no universally agreed upon normative ethical theory (like there is a universally agreed upon chemical theory of water) assume for a moment that normative ethical theorizing turns out to converge on a utilitarian account of goodness. Now consider twin-earth where everything else is the same as on earth except that the referent of the term 'good' is something other than maximal happiness. Do earthlings and twin-earthlings mean different things by their use of the term 'good' or are they disagreeing about the

nature of the property that ‘good’ refers to? Timmons suggests that here unlike the case of ‘Nixon’ or ‘water’ ordinary semantic intuitions underwrite the second rather than the first explanation.¹⁰ In any case, they are much less clear than in the typical cases of rigid designators, and this means that a posteriori identification of ethical properties with nonethical properties is much less plausible than in the typical cases of such identities.¹¹

Now, maybe Lycan would respond to this argument by suggesting that the reasons the intuitions are less clear is precisely because we have no widely accepted normative ethical account of goodness; as soon as we do, intuitions may come into line. Perhaps, but in my opinion, both sides of this argument depend on highly speculative views about the future development of ethical theorizing and the influence this will have on ordinary semantic intuitions. But notice that even if it turned out that ethical terms are rigid designators as Lycan suggests, there is still reason to be dissatisfied with the a posteriori identification as part of an account of the meaning of ethical claims. For this identity is a metaphysical identity that explicitly does *not* imply a semantic identity. Because of this, even if ethical facts can be shown to be identical to some set of natural facts in the way Lycan describes, Moore’s view that ethical claims cannot be synonymy-analyzed in nonethical terms is untouched by the reductive naturalist’s position. We are searching for an account of the meaning of ethical claims, and the descriptivist approach is to give this account in terms of the facts described, but if it turns out

¹⁰ Perhaps the second explanation is also correct in the case of ‘water’. Rosenberg (1994) provides an extended case against the Putnam-Kripke arguments from “twin-earth” intuitions. However, the point here is that, even if we grant the arguments in the case of ‘water’, there seems to be much less reason to grant them in Timmons’ “moral twin-earth” cases—for whatever force the intuition of difference of meaning is in the case of ‘water-twater’, it is much weaker in the case of ‘good-tgood’.

¹¹ This argument doesn’t undermine Sayre-McCord’s (1997) contention that ethical terms rigidly designate *moral* kinds. He argues that it is normative theory rather than scientific theory that determines the actual referents of our ethical terms, and that we would see twin-earthers as using terms with different meanings than our ethical terms, were we to discover (or stipulate) that their terms are not sensitive to the strictures of a broadly normative-justificatory theory. However, because he replaces natural kinds with moral kinds in the general kinds-semantics, Sayre-McCord’s view does not underwrite any reduction of ethical properties/facts to natural properties/facts (nor is it incompatible with such a reduction pursued in some other way).

that ethical claims and some set of nonethical claims describe the same set of facts without meaning the same thing, then descriptivism is an inadequate account of the meaning of ethical claims. So, in terms of the overall project in this part of the dissertation, descriptivism combined with a metaphysical reductivism that doesn't implicate a semantic reduction is not an adequate account of the meaning of ethical claims. Notice that this point applies also to those who may want to remain neutral on whether ethical terms are rigid designators, but nonetheless hold out hope for a type-type identity of ethical properties with natural properties along the lines of materialist reductions in the philosophy of mind or inter-theoretic reductions (e.g. temperature to mean molecular kinetic energy) in the sciences. Thus, since we are exploring the advantages of descriptivist accounts in this chapter, we can legitimately set aside the a posteriori identification position, for it simply doesn't provide a descriptivist account of the meaning of ethical claims at all.¹²

That would be the end of the story for descriptivist naturalism were it not for the fact that post-Smart developments in the philosophy of mind have provided a new *nonreductive* way to argue for materialism, and with it a new model for metaethical naturalism that is nonreductive and thus potentially immune to the objection just pressed against attempts to identify ethical facts with other prosaically natural facts. The new way to materialism in the philosophy of mind can be seen by considering an objection to type-type identification of mental and bio-physical properties. Many have objected that such identification has to be wrong, since it seems possible that there could be a being who has mental states without having anything like our bio-physical constitution. Think of a Martian or a robot feeling pain—maybe such thoughts are incoherent, but it's not obvious that they are. And until we can show that they are, it seems hasty to assume that the bio-physical properties cognitive science may come to correlate with mental properties

¹² To be fair to Lycan, I should mention that it is not clear that he intends to give an account of the meaning of ethical claims—at least not in (1988: ch 12).

should be viewed as being identical to these properties. Nevertheless, this need not undermine materialism. Functionalists (e.g. Sellars 1953, 1956 and Putnam 1960, 1963) have argued that we can characterize mental states by three kinds of causal relations: what causes them, what effects they have on other mental states, and what effects they have on behavior. This constitutes their “functional role”, and even if it might be realizable by radically different bio-physical properties that doesn’t undermine viewing mental properties as a sort of “higher-order” material property that gets instantiated solely by “lower-order” material properties. Of course, we cannot prove that there is no nonmaterial way to instantiate the mental properties, but cognitive science may be able to show that what fills these functional roles in us is bio-physical properties. This is no longer a strict type-type identity of mental and bio-physical properties, but it does provide a way to think of our mental properties as bio-physical in that they are functional roles solely instantiated by bio-physical properties. (Compare the way social properties are not clearly type-type reducible to biological properties but might nonetheless be thought to be functional roles which, for all of the societies we’re aware of, are solely instantiated by biological properties.)

This provides a new model for defending a posteriori naturalism: perhaps moral properties are functional roles that, at least in the cases we’re aware of, get instantiated solely by prosaically natural properties. In this sense, they wouldn’t be type-type identical to prosaically natural properties, but they might still be reasonably viewed as natural. This point is often put in terms of the *supervenience* of the moral on the physical. For instance, Brink, a proponent of such a “non-reductive a posteriori naturalism”, writes, “There is nothing strange and certainly nothing unique about the supervenience of moral properties on physical properties. Assuming materialism is true, mental states supervene on physical states, yet few think that mental states are metaphysically queer...”(1984: 120).

But supervenience is actually beside the point. What nonreductive a posteriori naturalists need is an argument for thinking that the ethical properties—which both sides of the debate

countenance and conceive of as supervenient on but not identical to natural properties—are themselves natural. And this, in turn, seems to require coming up with an account of what it is for a property to count as natural. Because this debate takes place on the backdrop of Mackie’s argument from queerness, typically naturalists adopt what might be seen as a conservative conception of the natural, in order to not beg the question against Mackie.¹³ That is, they typically characterize the natural in terms of our broadly scientific way of explaining, predicting, and coming to know the world, roughly: a property is natural iff it is investigable by science. And here the causal theory of reference can provide renewed hope for naturalism. Above, we saw that the problem with reductive a posteriori naturalism was that even if ethical terms turn out to rigidly designate prosaically natural properties, in that their use is causally regulated by these properties as many philosophers think that the use of ‘water’ is causally regulated by H₂O, identifying ethical properties with natural properties spells death for descriptivist naturalism. However, it’s now possible to treat ethical terms as rigidly designating *ethical* properties that are not identical to but supervenient on and instantiated by prosaically natural properties. That is, it is possible to develop a causal theory of the reference of ethical terms and thus to treat ethical terms as causally regulated by ethical *cum* physical properties. If this causal theory of the referents for ethical terms is correct, then ethical properties are, in principle, investigable by science; and so the ethical facts in which these properties figure are not identical to facts describable in nonethical terms, but we can still view ethical claims as gaining their meaning from describing a realm of natural facts.

¹³ Ridge (2003) offers a helpful list of different proposed characterizations of natural properties as “properties that (i) are the subject matter of the natural sciences (Moore 1903: 40), (ii) are invoked in scientific explanations (Little 1994: 226), (iii) would be identified by the best scientific theory and can be described in conceptual terms available to a being occupying a non-local point of view on the world (Crisp 1996: 117), (iv) can be known only a posteriori (Copp 2003, Moore 1903: 39), (v) can exist by themselves in time (Moore 1903: 41), (vi) confer causal powers (Lewis 1983), (vii) figure in the laws of nature (Vallentyne 1998), or (viii) explains similarity relations (Lewis 1983).” One might add (ix) that can be instantiated in space and time. (i) is obviously too conservative—it doesn’t allow for scientific progress—but by suggesting that the naturalist needs to adopt a conservative conception of naturalness, I mean to suggest that they should avoid characterizations like (iv), (v), (viii), or (ix). For all of these risk counting too much—or, more to the point, things that might invite the charge of queerness—as natural.

This is effectively the position Boyd (1985) defends. As part of a general causal theory of reference and knowledge for natural properties, he proposes that,

Roughly, and for nondegenerate cases, a term *t* refers to a kind (property, relation, etc.) *k* just in case there exist causal mechanisms whose tendency is to bring it about, over time, that what is predicated of the term *t* will be approximately true of *k*...Such mechanisms will typically include the existence of procedures which are approximately accurate for recognizing members or instances of *k* (at least for easy cases) and which relevantly govern the use of *t*, the social transmission of certain relevantly approximately true beliefs regarding *k*, formulated as claims about *t*..., a pattern of deference to experts on *k* with respect to the use of *t*, etc...When relations of this sort obtain, we may think of the properties of *k* as regulating the use of *t* (via such causal relations), and we may think of that as said using *t* as providing us with socially coordinated *epistemic access* to *k*...(1985: 195).

However, he rejects the view that in order for a property *k* to be viewed as natural, it must be semantically *or otherwise* type-type identifiable with physical properties. This is because, “there are natural kinds, properties, etc. whose natural definitions involve a kind of property cluster *together with* an associated indeterminacy in extension. Both the property-cluster form of such definitions and the associated indeterminacy are dictated by the scientific task of employing categories which correspond to inductively and explanatorily relevant causal structures”(Ibid.: 196). He offers ‘healthy’ as an example of a term that is causally regulated by a property-cluster, in this sense. And because type-type identities between higher-level properties and physical properties can be replaced by correlations between higher-level properties and physical property-clusters, he insists that “the moral realist can deny that it follows from the premise [that goodness is a natural property] that goodness is a physical property or that goodness has any physical definition, analytic or otherwise”(Ibid.: 199). This is because the realist can treat ethical terms as causally regulated by the ethical properties which are then instantiated by (clusters of) physical properties in much the same way that a materialist functionalist in the philosophy of mind might treat mental terms as causally regulated by mental properties which are then instantiated by (clusters of) bio-physical properties.

So Boyd defends a view according to which ethical terms are causally regulated by homeostatic property clusters, having to do with things like pain, pleasure, biological needs, psychological needs, social relationships, and complex interrelations among these things. This is the position he calls “Homeostatic Consequentialism”, and since he wants “the analogy between moral inquiry and scientific inquiry to be taken *very* seriously”(Ibid.: 204), he ultimately paints a picture of inquiry into the nature of goodness that is very similar to the way cognitive scientists might think of inquiry into the nature of pain or medical doctors might think of inquiry into the nature of healthiness. In all of these cases, references to the relevant property starts with ordinary language, but then investigation into its nature proceeds along the lines of pursuit of broad reflective-equilibrium that is characteristic of science in general.

The important point in all of this is that investigation of ethical properties is seen to be exactly parallel to the investigation of other “higher-level” natural properties such as pain or health that are not type-type identical to but nonetheless supervenient on and instantiated by physical properties. The fact that none of these “higher-level” properties is claimed to be type-type identical to physical properties saves this nonreductive a posteriori naturalist account from the object pressed against reductive a posteriori naturalism above, while the fact that ethical properties are claimed to causally regulate the use of ethical terms ostensibly reveals them to be scientifically investigable and thus, on the conservative conception of ‘natural property’, natural.¹⁴

Many have worried that this approach implicates a rather odd picture of moral theorizing, according to which, rather than considering things such as hypothetical cases and the consistency of various plausible moral principles, we should be empirically investigating what sorts of physical property-clusters causally regulate the actual use of our moral terms. Besides

¹⁴ A similar position is defended by Railton (1986).

being a highly revisionary conception of proper moral theorizing, this conception seems to have the undesirable result that, for any given range of investigated cases, the actual use of our moral terms in the range could be causally regulated in a mistaken or idiosyncratic way for which we would actually need an antecedently articulated moral theory to recognize and correct. Here the nonreductive a posteriori naturalist can seem to be in a bind. For in defending the naturalness of ethical properties, he has drawn a very tight analogy to other natural properties investigable by science; however, precisely because this analogy is so tight it seems that he's forced to a radically revisionary, if not downright immoral, conception of moral theorizing.

Brink (2001) attempts to overcome this problem by widening the conception of what it takes for an ethical term to be "causally regulated" by an ethical property and thereby what sorts of theorizing might be appropriate for investigating the nature of ethical properties. He suggests that in addition to the requirements on regulation articulated by Boyd, "any interpretation of regulation must also understand it in a way that is sensitive to various demands of consistency. To apply moral terms in virtue of N-properties is just to have a set of beliefs, however implicitly, about which N-properties are morally relevant"(2001: 168), but part of meeting the demands of consistency here is that these beliefs be revisable in light of consideration of counterfactual cases. Thus, ultimately Brink's view is that "a natural property N causally regulates a speaker's use of moral term 'M' just in case his use of 'M' would be dependent on his belief that something is N, were his beliefs in dialectical equilibrium"(Ibid.: 169), where the idea of "dialectical equilibrium" is supposed to capture the way in which good ethical theorizing proceeds in light of consideration of hypothetical cases and consistency rather than, as Boyd seems to think, by means of empirical investigation of the causal regulation of our actual use of ethical terms.

However, even if my moral beliefs are in dialectical equilibrium and thus my moral terms are regulated by some complex natural property cluster in just the way Brink claims, one might still worry that this doesn't vindicate the naturalness of ethical facts. For notice that an astrologist's

use of terms for the various astrological signs and their relations might be “causally regulated” in the dialectical way that Brink describes by real astronomical properties. But that doesn’t vindicate the naturalness of the putative astrological properties. For although in the envisaged case the “nonreductive a posteriori astrological naturalist” could claim that astrological facts supervene on and are instantiated by astronomical facts, unless she is claiming that the astrological facts are identical to the astronomical facts, we’d want to deny the existence of the astrological facts. Why? Probably because they seem to have some extra oomph or significance for those who believe in them that the mere astronomical facts could not have.

I think a parallel worry faces nonreductive a posteriori naturalism, which we can see by recalling that even the arch ethical *non*naturalist, Moore, thought that a thing’s ‘intrinsic value’ depended on its ‘intrinsic nature’, which he conceived of as a set of natural properties (Moore 1993, see Baldwin 2004 for more discussion of Moore on what we now call supervenience). But the supervenience of value on natural properties is arguably not enough to vindicate the naturalness of ethical facts. Why not? Some metaethicists (e.g. Gibbard 2003) have suggested that even though the open-question argument is fallacious, what Moore was onto in using it and ultimately denying the naturalness of ethical facts was their putative action-guiding character. I think there are two separate ideas packed into the intuition that putative ethical facts would be intrinsically action-guiding.

First, there is the *rationalist* idea that if it’s true that ϕ -ing in circumstances C is good, then anyone in C has a reason to ϕ . Now, it’s important to notice that in the phrase “reasons for action” hides an important distinction between what are sometimes called “motivating reasons” and “justifying reasons”.¹⁵ A motivating reason, or for short, a motivation is what *psycho-causally* explains why an agent acted in a particular way. There is a prominent debate here between Humeans who hold that motivation consists essentially of a belief-desire pair

¹⁵ I take this distinction from Smith (1994).

(Davidson 1978, Smith 1987, 1994) and anti-Humeans who hold that one can in principle be motivated without having a desire (Nagel 1970, McDowell 1978, Platts 1991, Wiggins 1991, Scanlon 1998, Shafer-Landau 2003).¹⁶ However, this is independent of debate about what constitutes a justifying reason, which is whatever normatively explains why an agent should act in a particular way—for obviously the fact that I believe ϕ is desirable or want ψ and believe that ϕ -ing will satisfy this desire will not always *justify* my ϕ -ing. The rationalist idea is about justifying reasons. For it is clearly false that the judgment that ϕ -ing is good in C entails that anyone in C will be appropriately motivated; after all, they might not recognize that ϕ -ing is good or even think it is bad.

Second, there is the *internalist* idea that one who judges that ϕ -ing in C is good will be ipso facto motivated to ϕ when in C. Here, it's important to notice that this thesis must be weakened to be plausible. Certainly we aren't *completely* motivated to follow our ethical judgments—otherwise there couldn't be weakness of will. Moreover, it seems wrong to say that the internal connection between ethical judgments and motivation is psychologically necessary. We can imagine cases where someone is capable of making an ethical judgment but is too motivationally overloaded to be motivated to do what she takes to be good or too depressed and listless to be motivated at all. This suggests that the thesis should be put in terms of fallible dispositions—viz., one who judges that ϕ -ing in C is good will be ipso facto fallibly disposed to be motivated to ϕ when in C. However, some externalists and even some internalists have noted that even this seems too strong. For there appear to be psychopaths who can judge that ϕ -ing in C is good but are only motivated to ϕ when in C if there is some external threat of reward or punishment. I think it's debatable whether psychopaths should be credited with making genuine ethical judgments rather than merely using ethical terms in a scare-quotes sense. However, Blackburn (1998: 61) presents the case of Milton's Satan who avows “Evil be thou my Good”, which is a

¹⁶ Alternatively, this debate is sometimes construed as about whether there can be “*besires*”—i.e. intrinsically motivational beliefs.

possibility that seems to make sense and “the fact that it makes sense casts some doubt on the very close identification...between ethics and practical, motivating states of mind”(Ibid.) But Blackburn presents an internalist analysis of these cases:

My own judgment on this debate is that externalists can win individual battles. They can certainly point to possible psychologies about which the right thing to say is that the agent knows what it is good or right to do, and then...does something else. And they can point to psychologies like that of Satan, in which it can be come a [motivating] reason for doing something precisely that it is known to be evil. But internalist win the war for all that, in the sense that these cases are necessarily parasitic, and hat they are parasitic upon is a background connection between ethics and motivation. They are cases in which things are out of joint, but the fact of a joint being out presupposes a normal or typical state in which it is not out.(Ibid.)

And this points the way to an even weaker statement of internalism—viz., insofar as one is normal, one who judges that ϕ -ing in C is good will be ipso facto fallibly disposed to be motivated to ϕ when in C.

In the context of universalist descriptivist accounts of the meaning of ethical claims, the rationalist and internalist ideas packed into the intuitively action-guiding character of ethics can be transmuted into claims about ethical facts. If rationalism and universalist descriptivism are true, then ethical facts are such that they are, as I shall put it, *intrinsically reason-giving*—i.e. the fact that ϕ -ing in C is good is a reason to ϕ in C. And if (weak) internalism and universalist descriptivism are true, then ethical facts are such that they, as I shall put it, *magnetize the will*—i.e. when one recognizes the fact that ϕ -ing in C is good, one is ipso facto, insofar as one is normal, fallibly disposed to be motivated to ϕ when in C. So, although both Moore and the naturalists can agree that two actions with all of the same natural properties cannot have different ethical properties (supervenience), Moore will insist that the mere fact that something has some set of natural properties is not enough to be intrinsically reason-giving or to magnetize the will, while the fact that something is good is enough to be intrinsically reason-giving and to magnetize the will. We might reasonably take this to be a reflection of the extra, so to speak, normative oomph of ethical judgments, which indicates that the supervenience (even in the style

of Boyd or Brink) of ethical facts on prosaically natural facts is not enough to vindicate the naturalness of ethical facts any more than the (hypothetical) supervenience of (putative) astrological facts on prosaically natural facts is enough to vindicate the naturalness of astrological facts. And it was probably this normative oomph of ethical facts that lead Mackie to charge them with being too queer to deserve a place in our ontology.

In the case of mental facts (or health facts) supervening on bio-physical facts, neither the higher or lower level sort of fact is thought to be intrinsically reason-giving or to magnetize the will; but in the case of ethical facts putatively supervening on natural facts, the two sorts of facts seem to have a different normative oomph. Thus, it looks like the nonreductive a posteriori naturalist needs to either deny that there is this difference or say more about why it doesn't undermine the claim that ethical facts are natural. Typically nonreductive a posteriori naturalists respond to this worry by denying the disanalogy. That is to say, they deny that ethical facts have the normative oomph or action-guiding character that Moore seemed, however inchoately, to be onto with his open-question argument.¹⁷

But we should recognize the costs of this move. If the naturalist denies the action-guiding character of ethical facts, doesn't he violate two very intuitive ideas—viz., rationalism and (weak) internalism? Naturalists such as Boyd and Brink don't see this as a problem because they take themselves to have good reason for denying rationalism and internalism. And this underwrites their response to Mackie's argument from queerness. They deny premise (2) by rejecting rationalism and internalism. Although this has the obvious advantage of offering a response to Mackie, I think it is also what generates the most problems for nonreductive a posteriori naturalists. To draw out the disadvantages of the view, it will prove useful to look at Boyd's

¹⁷ Another possible response would be to admit the disanalogy but deny that it undermines the claim that ethical facts are natural by arguing that the extra normative oomph of ethical facts doesn't undermine their naturalness. I doubt this will work unless we have some reason to think that other facts with this sort of oomph count as natural, but it's not even clear that there other facts with this oomph, let alone that they are natural.

argument against rationalism and Brink's argument against internalism; in both cases, I think they miss the point of these intuitive doctrines.

Boyd rightly notices that "it has often been objected against moral realism that there is some sort of logical connection between moral judgments and reasons for action which a moral realist cannot account for"(1988: 214). And he admits that, "It is of course true that the naturalistic moral realist must deny that moral judgments necessarily provide reasons for action...it is hard to see how the naturalistic moral realist can escape the conclusion that it would be *logically possible* for there to be a human being for whom moral judgments provide no reasons for action"(Ibid.). His response to this objection is to endorse what he calls the "standard naturalist response" of explaining "the apparent intimacy of the connection by arguing that the natural property moral goodness is one such that for psychologically normal humans, the fact that one of two choices is morally preferable will in fact provide some reason for preferring it"(Ibid.: 215)¹⁸ So the core idea in this response to the objection from rationalism is that we don't need rationalism to explain the intuitively intimate connection between moral judgment and reasons for action because moral facts hook into things that normal humans already prefer.

This response is effectively in line with Mill's (1861) response to the objection that judgments about which actions are good—conceived of as which actions maximize utility—are not intrinsically reason giving. Mill admits that they aren't intrinsically reason giving, but he argues that they are reason giving for those who have been appropriately inculcated into human

¹⁸ To be fair, I should mention that he says that this response doesn't completely rebut the objection because it ignores "the very strong intuition which many philosophers share that the person for whom moral judgments are motivationally indifferent would not only be psychologically atypical but would have some sort of *cognitive* deficit with respect to moral reasoning as well"(Ibid.). However, as far as I can tell, Boyd runs consideration of rationalism and internalism together here. Whether one is appropriately motivated by one's moral judgments is a different issue from whether moral judgments commit one to reasons to act. Of course, by endorsing a very strong form of reasons internalism, we can construe the former as a necessary condition for the latter; but I see neither a reason to endorse such a strong form of reasons internalism nor why doing so will help Boyd out of the objection I go on to discuss in the text above.

society because such people will have had their preferences shaped so that they prefer (at least to some degree) maximizing utility. Given how close Boyd's answer is to Mill's, it's unfortunate that Boyd doesn't address either of the two standard rejoinders to this response to the objection from rationalism.

First, while normal and appropriately socially inculcated humans may prefer the well-being of people other than themselves, it seems to be a plainly false that people prefer the "greater good", indeed someone who prefers maximizing utility *per se* is psychologically abnormal and maybe even ethically abhorrent.¹⁹ And Boyd's more nuanced "homeostatic consequentialism" doesn't escape this worry. For, according to his proposal, "the homeostatic cluster which defines moral goodness...[is] social rather than individual. The properties in homeostasis are to be thought of as instances of the satisfaction of particular human needs among people generally, rather than within the life of a single individual"(Ibid.: 204fn). But if this is the case, then it simply seems false that psychologically normal humans will necessarily prefer what is morally good. Of course, they will prefer that their own needs and the needs of their family-social group be satisfied, but satisfaction of these needs alone doesn't constitute the homeostatic property cluster that Boyd thinks causally regulates the use of the term 'good'.

Second, even psychologically normal people *did* prefer the greater good or the satisfaction of human needs among people generally, it requires an argument to show that this preference and the fact that a particular action will satisfy it gives one a reason to perform the action. For although a preference and the fact that some action will satisfy it (when recognized) may be the motivating reason that explains someone's action, that does not imply that one has/had a justifying reason to perform the action. But the rationalist idea is that judging that an action is good (or ought to be performed) is ipso facto to be committed to there being a justifying reason to perform it. Boyd might want to deny the "ipso facto", but if so he needs to replace it with

¹⁹ For a similar point see Wolf (1982).

something that maintains the strong tie, and I cannot see that he has. As such, the nonreductive a posteriori naturalist is stuck denying not just the letter of but even the intuition behind rationalism, which is a serious disadvantage.

Brink (1984: 42) runs together what I've been calling rationalism and what I've been calling internalism under the label "internalism", which is apparent from his characterization of internalism as consisting of at least three distinguishable components. The first is the thesis that "moral claims *necessarily* motivate or provide a reason for action". The second is the thesis that the first thesis follows from the concept of morality and so is an a priori claim; he writes, "Since it is the *concept* of morality that shows that moral considerations necessarily motivate or provide reasons for action, this claim about the motivational power or rationality of morality must be a priori"(Ibid.). The third is that "since it is the *concept of morality* that determines this fact, the rationality or motivational power of moral considerations cannot depend on any substantive considerations such as what the content of morality turns out to be, facts about agents, or the content of the correct theory of rationality"(Ibid.). He admits that internalism, so construed, represents "a very narrow and strong claim," but he suggests that "it is this claim that underlies and explains traditional internalist claims," and, in particular, the internalism behind Mackie's argument from queerness ((Ibid.).

I think he is right that Mackie has both the reason-giving and the magnetizing character of morality (perhaps inchoately) in mind when he charges moral facts with being queer, but I seriously doubt that this strong triad of theses is necessary to underwrite widespread intuitions about the action-guiding character of morality. But let us set aside rationalism and just focus on internalism. To see that Brink's characterization is too strong, consider the second and third members of the triad. Is this family of views a priori and not dependent on substantive moral issues because it follows from the concept of *morality* as Brink suggests? In my view, the family of views might be a priori and nonsubstantive, but this needn't be because it follows from the

concept of *morality*. All Mackie needs from internalism (and, in my view, all he intended) is that the deployment of specific ethical concepts such as ‘good’ and ‘ought’ in making ethical judgments is appropriately conceptually connected to the motivational dispositions of the person making the judgment. In this sense, internalism is a doctrine about particular ethical judgments—and so might be seen as conceptual claims about the particular ethical concepts deployed therein—and need not be about the concept of morality *per se*.

Moreover, as we saw above, this connection needn’t be one that makes motivation a necessary consequence/component of ethical judgment. Rather, it is much more plausible to think that the connection is to fallible dispositions to be motivated and contingent on the agent being normal. The idea behind such a weak form of internalism is that if someone who appears to make an ethical claim isn’t appropriately disposed to be motivated, we will think she is abnormal, insincere, or that she hasn’t deployed particular ethical concepts at all. This is still a controversial claim, but whether or not it is a priori and independent of substantive considerations in a theory of morality or rationality will depend on one’s views of the nature of a priori knowledge and the strictures of theory building that apply to these theories. It seems to me that whether the connection between moral considerations and motivations claimed by the internalist is true is indeed a conceptual issue, but that doesn’t mean that it is an issue that each of us can decide individually by a priori reflection on our concept of morality—rather we need a careful account of the nature of our collective particular moral concepts such as ‘good’ and ‘ought’; and this may be precisely what substantive moral theorizing is trying to do.

Brink offers many reasons for rejecting various stronger forms of internalism, but the closest he comes to attacking something as weak as the weak form of internalism discussed above is citing the putative counterexample of the “amoralist” who makes ethical judgments but is completely unmotivated to act in accordance with them. Since the amoralist is imaginable, Brink argues that the possibility of his existence cannot be ruled out by our ethical concepts. Of course,

the way we've defined (weak) internalism, as long as the amoralist is thought to be abnormal, he doesn't fall under the extension of the conceptual claim, and thus doesn't really pose a problem. But I imagine that Brink will respond by saying that we could also imagine a whole world of amoralists in which the disconnect between their moral claims and motivation could not be thought to be abnormal.

Does this undermine (weak) internalism? I think not. There seem to be two ways to defend the thesis that there is a conceptual connection between making ethical judgments and being appropriately disposed to be motivated to act. First, one could argue that Brink is mixing up modalities: Even if individual a priori reflection cannot rule out the possibility of someone who makes moral judgments but is not at all motivated (or doesn't at all have reasons to be motivated) to act in accordance with them, that doesn't mean that the amoralist is, so to speak, morally possible. The best theory of the nature of moral concepts might say that one who is not at all motivated to act in accordance with a set of apparent moral judgments isn't really making moral judgments at all, in which case, although the amoralist is in some sense imaginable, she is not, given the theory, possible. Second and perhaps more plausibly, one could argue that in considering the world full of amoralists, it actually seems that they aren't deploying ethical concepts, and so they represent no counterexample to the internalist point about our ethical concepts.

I admit that intuitions run thin here and so appeals to intuitions on both sides of the debate can seem to be cooking the books, but I think that nonreductive a posteriori naturalists tend to attack a straw man version of internalism in response to the worry that their view is incompatible with internalism. As we've seen, weaker versions of internalism are very plausible. But, weak as they are, even if they are true, they pose a challenge for a posteriori naturalism: we want to know why ethical facts as the naturalist construes them have the normative oomph rationalism and internalism would force them to have? This extra oomph ethical facts over and

above the natural facts on which they putatively supervene seems to either undermine the supervenience claim or call into question the appeal to this supervenience claim in order to vindicate the naturalness of ethical facts.

But maybe this vindication isn't as hard as I've made it out to be. Recall the dilemma for the nonreductive naturalist that got us to this point: He tried to vindicate the naturalness of ethical facts by arguing that they supervene on natural facts in the same way that mental facts supervene on physical facts, but the action-guiding character of ethical facts seems to belie any tight analogy. And since it is precisely this character that leads to Mackie's argument from queerness, the nonreductive naturalist was forced to either deny the action-guiding character of ethical facts or to explain why this character doesn't undermine their claim to naturalness. However, we haven't yet explored the second horn of this dilemma. Certainly the fact that there is an internal connection between ethical claims and reasons for and motivation to action doesn't *show* that putative ethical facts would have to be nonnatural. For all parties to this debate would agree that the pain is natural, but it also seems that there is an intrinsic connection between claims about painfulness and reasons for and motivation to action. Indeed, this connection seems to be much like the internalist connection we reached in our discussion of (weak) internalism. Internalism about pain-judgments would say that one who judges that something is painful is, if normal, ipso facto fallibly disposed to be motivated to avoid what he judges to be painful. And rationalism about pain-judgments would say that if it is true that something is painful then one has reason to avoid it. These are plausible internalist and rationalist theses about pain judgments, but if so, then it isn't clear why their plausibility in the ethical case causes trouble for the metaethical naturalist. Can't he say that ethical facts are natural and action-guiding just like facts about pain?

He can, but this would be false. Pain is itself a motivational psychological state, which can be explained in terms of physical and psychological damage that it is evolutionarily advantageous

to be disposed to avoid. So there is a very easy explanation for why one who judges that something is painful will usually be motivated and has reason to avoid it. The situation in the ethical case is different. Unless one wants to drop into a subjectivist view about ethical judgments, goodness is not itself conceived of as a motivational psychological state. It's more objective, and this is what seemed to worry Mackie in the first place: it's objective but also intrinsically reason-giving and magnetizing of the will. If we transposed Mackie's argument from queerness into an argument about painfulness judgments rather than ethical judgments, the first premise would be much less forceful—painfulness judgments are precisely dependent on subjective reactions in a way that it seems that ethical judgments are not. So, I think the (appropriately weak) sense in which ostensible ethical facts are action-guiding continues to pose a special sort of problem for vindicating their naturalness by appeal to supervenience *a la* nonreductive a posteriori naturalism.

In this section, I have explored one of the most prominent attempts to naturalize ethical facts in response to Mackie's argument from queerness. I think the attempt is initially well motivated. Mackie's conclusion that affirmative ethical claims are uniformly false should be avoided, but so too should the posit of mysterious entities. A typical way to show that some entity is nonmysterious is to show how it can be seen as natural and thus in the purview of the natural sciences. What we've seen here, however, is that a priori *identification* of ethical facts to natural facts suffers at the hands of something like Moore's open-question argument. Reductive a posteriori naturalism modeled on Smart's reduction of mental properties to bio-physical properties or the Kripke-Putnam reduction of water-properties to H₂O properties also has problems that seem avoidable only at the cost of giving up the descriptivist approach to the semantics of ethical claims—for we seemed forced to agree that even if the moral and natural facts are identical, the meanings of the claims are different. And nonreductive a posteriori naturalism turns out to be incompatible with a plausibly weak versions of rationalism and naturalism. In the end, it seems to me that even if a posteriori naturalists can identify some facts

that are both natural and causally regulate our ethical claims, they haven't succeeded in showing these facts to be the facts described by our ethical claims. In the following section, I shall explore another attempt to show that ethical facts need not be mysterious, which provides an alternative response to Mackie's argument from queerness.

II. NONNATURALISM

If one thinks that the meaning of ethical claims is to be explained in terms of their describing a realm of ethical facts that are universal in the sense of being objective and independent of variable human practices, but one despairs at the prospects of showing intrinsically action-guiding ethical facts to be natural, one might argue that the facts ostensibly described by ethical claims are *sui generis* and then argue that despite this and *pace* Mackie they should be admitted into our ontology. In contemporary metaethical debate, this view was first articulated, defended, and made popular by Moore (1903). Moore thought that other ethical concepts could be semantically analyzed in terms of the concept of 'good', which he claims to be "the only simple object of thought peculiar to Ethics"(Ibid.: 5); and he thought this analysis proceeds in a recognizably consequentialist fashion: "To assert that a certain line of conduct is, at a given time, absolutely right or obligatory is obviously to assert that more good or less evil will exist in the world, if it be adopted than if anything else be done instead"(Ibid.: 25). And although he is famous for arguing that 'good' is unanalyzable, he says that with this term he means, the "unique object—the unique property" we have before our minds when we say that something has "intrinsic value," "intrinsic worth," or "that a thing ought to exist"(Ibid.: 15). Nevertheless, Moore clearly thinks that 'good' is the most basic ethical concept. But rather than following previous consequentialists in giving a reductive analysis of what goodness is; Moore argued that the term 'good' is indefinable and that goodness is therefore a simple nonnatural property.

The centerpieces of his defense of this view was his identification of the so-called “naturalistic fallacy” and, what we’ve already seen briefly above, his deployment of the “open-question argument”. According to Moore, one committed the naturalistic fallacy when one sought to infer that something is (ethically) good solely from its nonethical properties.²⁰ One common way to try to justify such an inference was by appeal to an analysis of goodness in terms of some complex of nonethical properties; for, if we know that ‘good’ just means ‘*N*’, where ‘*N*’ refers to some (perhaps complex) nonethical property, then we can infer from the fact that something is *N* to the conclusion that it is good. Moore sought to undermine all such synonymy-analyses with his open-question argument.

As we’ve already seen, the basic idea behind the open question argument was to ask, for any proposed synonymy-analysis *N*, whether the question “Is *N* good?” is open or closed. There are several related things Moore seems to have meant by “open”: An open question is one that can be asked “with significance”; an open question is one where having the definiendum “before the mind” is not sufficient (in those competent with the concepts) for having the definiens also “before the mind”; an open question is one that can be the expression of a “serious doubt”(Ibid:15-17). In any case, according to Moore, if the question is open for those who are competent with the relevant concepts, then the analysis fails. For example, Moore considers the suggestion that goodness is that which is desired to be desired, but since the question: “Is that which is desired to be desired good?” is open, Moore concludes that this analysis must be incorrect. Moreover, he thought that all naturalistic analyses of goodness were bound to fail the test: it certainly seems that all extant analyses fail the test, and there is no good reason to suppose that the history of failed attempts is not indicative of the fact that no analysis can pass

²⁰ As many commentators have pointed out, Moore’s label “the naturalistic fallacy” is horribly misleading since Moore thinks that exactly the same mistake is made by ethical *super*-naturalists who seek to reduce ethical facts to facts about the wishes of divine beings, and the putative mistake is not really a *fallacy* but rather the conclusion of a potentially valid line of argument whose soundness Moore disputes. See Ridge (2003).

the test. This led him to conclude that ‘good’ is indefinable and thus goodness is a simple *sui generis* property.

There are several problems with this line of argument, but as I noted above it is widely agreed that Moore *was* successful in undermining synonymy-analyses of the concept ‘good’ in nonethical terms. And, more generally, even if one disagrees with Moore’s claim that other ethical claims can be analyzed in terms of maximizing good outcomes, it seems that his open question argument provides grounds for the view that *no* ethical claims can be synonymy-analyzed in nonethical terms. It was perhaps anachronistic on Moore’s part to read previous ethical theories as attempts to provide a semantic analysis of ethical claims; however, insofar as we’re after an adequate account of the meaning of ethical claims, Moore’s conclusion provides a significant restriction of the possible options.

Above we saw that a posteriori naturalists argue that they can grant Moore’s point while consistently pursuing an a posteriori reduction of ethical facts to prosaically natural facts. Against Moore’s position, that is, they object that that even if all synonymy-analyses of the *term* ‘good’ fail, it is highly implausible to suppose that this is probative with respect to the nature of the *property* referred by ‘good’ and, correlatively, with respect to the issue of whether ethical facts can be reduced to some set of nonethical facts. That is, Moore’s argument for nonnaturalism illicitly assumes that a linguistic definition of a term must simultaneously serve as an analysis of the referent of the term. But, as we’ve already seen with analyses of ‘water’ and ‘H₂O’, this assumption doesn’t generally hold. Given the post-positivistic distinction between concepts and properties, this is an obvious mistake in Moore’s argument. But I think one seriously underestimates the prospects for nonreductive accounts of ethical facts if one dismisses them because of this mistake in Moore’s argument. It is not clear to me that Moore even needs the open question argument to support the view that ethical facts are *sui generis*, since one might grant that an analysis’ failure to pass the open-question test doesn’t undermine

it and still think that Moore was right that all extant analyses of the property of goodness (including a posteriori analyses) in nonethical terms have turned out to be incorrect. That is, even if the semantic intuitions of those competent with the term ‘good’ aren’t probative with respect to the correctness of a proposed property-analysis, one might reasonably think that the ethical intuitions that fund much ethical theorizing *are* probative or at least highly suggestive. It’s in this vein that Shafer-Landau writes,

...there are two reasons to think that the open question argument does in fact create a substantial burden of proof against reductive naturalism. The first is that most continue to question the alleged moral-natural identities even after a very great deal of substantive investigation has been done...It may just be a case of having yet more work to do. These questions might indeed be closed...But we don’t, right now, have any reason to suppose that they questions that seem so clearly open are, unbeknownst to so many of us, in fact closed...And that is because, secondly, we have yet to see a defensible, well-developed naturalistic semantics that, after excellent application, yields the conclusion that ethical reductionism is our best option.(2003: 67)

As Shafer-Landau indicates, this is a burden-shifting move. Even Moore would have agreed that the open question argument does not *establish* that goodness cannot be analyzed in nonethical terms, since this would require establishing the nonexistence of a correct analysis in nonethical terms and Moore only “tests” (a few of) those proposals on offer. However, he took himself to have provided a presumptive case against all reductive analyses, and even if one rejects his way of arguing against the proposals on offer, the fact that no reductive analysis has emerged from centuries of ethical theorizing can likewise provide a presumptive case against reductivism.

Of course, as we’ve seen, these considerations do not imply that nonnaturalism is the only viable descriptivist universalist account of the meaning of ethical claims. The arguments from Moore and Shafer-Landau canvassed so far support only the conclusion that the property of goodness and correlatively ethical facts are *sui generis*, but nothing about these arguments rules out the possibility that ethical facts are nonetheless *natural*—and the view that they are is basically the nonreductive naturalist position that I discussed above. Of course, we saw that this view has problems, but if one finds nonnatural facts mysterious, one might prefer to accept

these problems rather than accepting the posit of nonnatural facts. In the face of this, there are two strategies one might pursue in defense of nonnaturalism.

First, some philosophers (e.g. Gibbard 2003) have suggested that what Moore was really onto in arguing for nonnaturalism was the truth of ethical internalism. As we've seen, a weak version of ethical internalism is highly plausible. But how could internalism be true if the facts described by ethical judgments are universal and objective in the sense of being independent of particular human practices, patterns of reasoning, etc.? Well, they have to be a special sort of fact, such that when one judges that one of them obtains, one is (conditional on being normal) (at least somewhat) motivated (or has reason to be motivated) to act as they dictate. On a typical, albeit controversial, view, natural facts aren't like this. When we empirically investigate the natural world, we don't seem to find facts that "magnetize the will". At least as far as recognition of natural facts goes, motivation seems to require not just recognition but an additional subjective engagement with the facts. So, if one wants to be a descriptivist universalist, it can seem that one needs a special realm of nonnatural facts that magnetize the will in order to explain the truth of internalism.

If one accepts that facts which of their own accord magnetize the will are not easily integrated into the modern conception of the natural world, then this line of argument provides some motivation for nonnaturalism over naturalism. But the argument is unlikely to move the rabid naturalist like Mackie: if one's sole reason for being a nonnaturalist is that it is the only way to hang onto descriptivism, universalism, and a weak form of ethical internalism, then, given the mysterious nature of putative nonnatural facts, this should be taken as a problem for the conjunction of those three views rather than a good argument for of nonnaturalism.

This has encouraged the development of a second strategy in favor of nonnaturalism over nonreductive a posteriori naturalism. Shafer-Landau (2003) and Regan (2003) have recently defended nonnaturalism on epistemological rather than moral psychologistic grounds. Shafer-

Landau argues that the key to understanding the debate between the nonnaturalist and the nonreductive naturalist is not ontological parsimony but epistemic continuity. Both nonnaturalists and nonreductive naturalists agree that there are prosaically natural facts about such things as tables and chairs, C-fiber firings, laws of nature, etc. Moreover, most naturalists are nonreductive naturalists, which means that they agree with nonnaturalists that ethical facts are not identical to some set of these prosaically natural facts, which means that *both* camps accept that there are *sui generis* ethical facts.²¹ What they disagree about is whether our epistemic access to ethical facts is continuous with our epistemic access to naturalistic facts. Shafer-Landau writes,

...now there seems very little difference, especially in matters of ontological parsimony, between nonreductive naturalism, and nonnaturalism. Both are kinds of property pluralism, Both can have identical ontological inventories. The only difference is in whether or not the theorists will view the laws that govern the properties contained therein as discoverable all in the same “scientific” way.(2003: 64)

There are perhaps various accounts of the “scientific” way in which we come to know natural facts, but, on this way of construing the debate, nonreductive naturalism is constituted by a thesis about this general epistemological issue and then a thesis to the effect that moral knowledge is acquirable in essentially the same way as knowledge of other natural facts; and nonnaturalism is likewise constituted by a thesis about the general epistemological issue and then a contrary thesis to the effect that moral knowledge is acquirable only in an essentially different way from knowledge of natural facts.

Of course, naturalists will respond by objecting that although ethical knowledge has a special practical import, this doesn’t mean that the way we come to know it is non-scientific. What can

²¹ To be more precise, we should note that there are different axes along which two facts can be *sui generis*. For, of course, naturalists will insist that both ethical facts and other kinds of fact are of the same generis: *natural*, while at the same time agreeing with the nonnaturalist that ethical facts are not, e.g., *physical* or *psychological*.

decide this issue? Shafer-Landau argues that the study of goodness—i.e. ethics—is not a scientific way of coming to know what the ethical facts are because ethics is concerned more with practical prescription and recommendation than prediction and causal explanation. And although this initially begs the question against naturalists, something seems clearly right about the claim that ethical knowledge is peculiarly practical. If a parent knows that he ought to care for his children, then this provides (perhaps overriding) grounds for engaging in particular sorts of conduct. Knowledge that water is H₂O or even that eating meat causes pain do not seem to have this same sort of practical import. But does this practical character of ethical knowledge imply that the way we come to have ethical knowledge is distinct from the way we come to know prosaically natural facts?²²

Shafer-Landau attempts to show that it does by arguing for what he calls of “moral reliabilism”, which is the view that verdictive ethical beliefs²³ “can be epistemically warranted, provided they emerge from a reliable belief-forming process”(Ibid.: 272). Here, he draws on Goldman’s version of process reliabilism for empirical knowledge, but he thinks that one can preserve the reliabilist idea without the usual “causalist” assumption that reliability is to be explained in terms of causal interaction between the known fact and the reliable process that produces beliefs about this fact. Rather, although he begs off the task of articulating particular reliable belief-forming mechanisms for ethical beliefs, he proposes a point from which such an account could plausibly start. According to him, a good place to look for reliable ethical-belief-forming processes is our reliance on moral exemplars. He writes, “by devoting careful attention to the doings of paradigmatically immoral individuals, we may try to identify some common flaws that undermine correct moral perception and deliberation. We naturally extend the

²² In the following three paragraphs, I draw on my (2005b).

²³ He defends a different epistemological view for knowledge of moral principles, according to which our knowledge of many moral principles is self-evident. I criticize this view in my (2006).

investigation so as to capture the processes of moral assessment undertaken by those we deem as most ethically sensitive and insightful”(Ibid. 296).

However, it is unclear why consideration of exemplars should, even in principle, help to identify *noncausal* reliable belief-forming processes. Is the idea supposed to be that they have a special quasi-perceptual access to the nonnatural facts? If so, Shafer-Landau’s moral reliabilism is just a disguised version of Moorean intuitionism, which we all know to be more of a promissory note than a genuine epistemology theory. However, Shafer-Landau’s rejection the causalist assumption is important because, as far as I can tell, it is the only thing that divides his moral epistemology from a would-be a posteriori naturalist *cum* moral reliabilist; and since his reason for being a nonnaturalist is that he thinks that the way that we come to know moral facts is importantly different from the way that we come to know natural facts, if there is no difference here, Shafer-Landau cannot claim to have provided an alternative to a posteriori naturalism.

Moreover, I think Shafer-Landau’s rejection the causalist assumption generates specific problems for his version of moral reliabilism. In the case of empirical knowledge, reliabilism is, so to speak, explanatorily closed, meaning that the general (roughly, causal-tracking *cum* a posteriori theory-building) template for explaining the epistemic status of garden-variety empirical beliefs itself applies to the special case of explaining the epistemic status of empirical beliefs about which belief-forming processes are reliable. And this template is highly plausible because it is a part of the general causal-explanatory model provided by our best (highly predictive and explanatorily successful) overall scientific theory of nature (including the place and location of human minds therein). Shafer-Landau’s template for explaining the epistemic status of garden-variety moral beliefs is vague, but (having denied the causalist assumption) this much is clear: it is not a species of the more general reliabilist template for explaining the epistemic status of empirical beliefs, nor is it explanatorily closed. Thus, the explanatory

template he offers doesn't inherit the plausibility of our best scientific theory and can't generate its own explanation of the positive epistemic status of believing in particular applications of it. As a result, I wonder what (if anything) is supposed to generate positive epistemic status for beliefs about which belief forming processes are reliable, when these processes are not causally related to the relevant facts.²⁴

But it should have seemed wrongheaded from the beginning to try to defend nonnaturalism by arguing that ethical knowledge is different from knowledge of natural facts and then to appeal to a reliabilist model of knowledge. For even if a difference between knowledge of ethical facts and knowledge of natural facts can be made out within the reliabilist conception of knowledge, this won't capture the distinctive practical nature of ethical knowledge. I think Regan (2003) provides a more promising sort of argument for the view that ethical knowledge is essentially different from knowledge of prosaically natural facts. This comes in his attempt to respond to Korsgaard's objection to Moorean "substantive realism": she finds it implausible to think that "we have normative concepts because we've spotted some normative entities, as it were wafting by" or that "ethics is really a theoretical or epistemological subject"(1996: 44). Regan argues that there is no necessary connection between thinking of ethical knowledge as knowledge of a special sort of fact which requires a special sort of epistemic access and thinking

²⁴ To be fair, since Shafer-Landau holds that moral facts supervene on natural facts, he could appeal to general (empirical) reliabilism to explain the positive epistemic status of true moral beliefs in terms of processes that reliably track these underlying natural facts. (Thanks to Michael Ridge for pointing this out to me.) Thus, the story about what makes moral belief forming processes reliable needn't be even mostly noncausal. However, even if moral facts supervene on natural facts, it's unclear that moral beliefs are always formed by processes that track these underlying facts, since the supervenience base may be too complex to track. Now, perhaps there are complex cognitive-social mechanisms that aid in tracking the supervenience base, or perhaps (as Shafer-Landau has suggested to me) there is a simpler "resultance base" that is easier to track. But if so, it's unclear that the resulting moral epistemology is distinct from what a reliabilist-naturalist would endorse; and since what's crucial to Shafer-Landau's *non*naturalism is that it offers a distinctive moral epistemology, his position threatens here to collapse into a form of naturalism. Or if it doesn't, this must be because reliably tracking the supervenience base or the resultance base of moral facts is insufficient for epistemic access to moral facts, but then we need some story about the link between the causally explainable reliable tracking of the natural facts that underlie moral facts and the not causally explainable reliable formation of moral beliefs. Here one is tempted to say that the link requires a special moral intuition—but of course that won't do.

that these facts are (made up of) entities “wafting by” that we come to know by observation and theoretical reasoning.

To this end, he describes a prominent conception of the (phylogenetic and ontogenetic) co-development of human agency and consciousness: We start with brute drives, but because of the competition of brute drives we develop a higher-order drive or mechanism for arbitrating between conflicting drives. Then, with the ability to consciously represent the goals of our drives, we develop desires and correlated higher-order desires. This becomes more and more complex and calibrated by higher- and higher-order desires.(2003: 652-654) According to Regan, however, this much is not enough for agency since with this much we are merely “mechanisms generating behavior...[and not yet] agents who choose and execute actions”(Ibid.). To become agents we must become capable of asking what Regan calls the “desire-transcending question”: “whether our highest-order desires are sensible, worth pursuing, justified—or, if we are procedurally more coherentist, whether our whole constellation of desires is sensible, worth pursuing, justified”(Ibid.).

Regan thinks we clearly do regard ourselves as agents and thus regard ourselves as being able to ask the desire-transcending question, but in regarding ourselves as such, we presuppose that there is some standard independent of our own desires by which the question, “What ought I do?” can be answered. This independent standard is goodness—or what Regan (confusingly) calls “the Good”²⁵: “if we are to regard ourselves as agents, we need the Good”(Ibid.: 654). Thus, according to Regan, our way of coming to know the Good is fundamentally different from the way in which we come to know prosaically natural facts about things such as what objects are yellow or what the chemical constitution of water is. Knowledge of goodness, in his view, comes from critical (transcendental) reflection on the operation of practical reasoning, while knowledge of natural properties comes from the observation and speculative reasoning. He

²⁵ It’s confusing because Moore uses “the Good” to refer to the naturalistic supervenience base of the nonnatural property goodness, but Regan clearly intends “the Good” to be something nonnatural.

writes, “‘Good’ is nonnatural because it is not accessible to speculative reason; it is a property (for practical purposes) because we are compelled to posit a standard for our will which is external to our inclinations, as truth is an external standard for our speculative beliefs”(Ibid.: 656).

So, on this view, we wouldn’t have ethical knowledge if we didn’t have practical agency, but we might nonetheless have much knowledge of prosaically natural facts. In this way, Regan thinks that knowledge of ethical facts is discontinuous with knowledge of prosaically natural facts. He writes,

The reason ‘good’ is neither descriptive nor expendable is that it signals a different way of looking at the world from the way signaled by ordinary use of descriptive property terms. It is one thing to look at the world and see what is there. It is another thing to consider possibilities for the future and recognize which of them should be realized. Descriptive terms are the stuff of the first project ‘good’ is the stuff of the second.(Ibid.: 659)

So what emerges from this discussion of Moorean nonnaturalism in its updated form due to Regan is a view of ethical facts as *sui generis* and known in a *sui generis* way through practical reasoning rather than speculative reasoning or observation. In contemporary metaethical discussions, nonnaturalism is often dismissed out of hand because it violates the naturalist dogma according to which the only facts are facts that can be discovered by empirical investigation and observation. I think Shafer-Landau is right to point out that (depending on how we count) the metaethical nonnaturalist’s ontological budget can be seen as parsimonious as that of metaethical naturalists like Brink and Boyd. Moreover, I think Regan provides a better argument for nonnaturalism than the mere appeal to the combination of descriptivism, universalism, and internalism: He argues that ethical knowledge is discontinuous with speculative knowledge of natural facts in that it plays an atypical practical role in our cognitive economy.²⁶ And, to my mind this reveals step (4) of Mackie’s argument from queerness to be

²⁶ This is, of course, a point that was hinted at a long ago by Frankena: “To my mind, what makes ethical judgments seem irreducible to natural or metaphysical judgments is their apparently normative

more problematic than it is usually thought to be. In effect, nonnaturalism can pit our commitment to our conception of ourselves as agents against our commitment to naturalism and then find the former to be stronger than the latter.

Nonetheless, I think the case for nonnaturalism is inconclusive. Given the distinction nonnaturalists want to draw between ethical knowledge and speculative knowledge of natural facts, it is unclear to me why we need to posit ethical facts at all. Consider the previous quote from Regan. Someone could agree that the concepts expressed by descriptive property terms and the concepts expressed by ‘good’ operate in exactly the way Regan claims and nonetheless go on to deny that ‘good’ describes a special ethical sort of property that figures in a special nonnatural sort of ethical facts. What, from the nonnaturalist’s perspective, would be wrong with the following picture: We have two fundamentally different ways of engaging with the world that line up with the possible conclusions or presuppositions of pieces of speculative and practical reasoning. The speculative way of engaging with the world involves making judgments and expressing them in claims that aim to describe objective facts; thus, these claims and judgments carry ontological commitment to the existence of the facts the purport to describe. But the practical way of engaging with the world involves making practical judgments and expressing them in claims that aim to regulate actions rather than to describe objective facts; thus, these claims and judgments carry no particular ontological commitments.²⁷

I think the problem nonnaturalists would have with this picture is that it makes the validity of ethical judgments ultimately relative, in the sense that they are answerable to the internal regulation of practical reasoning and not to some independent and objective ethical reality. In contrast, Moore seems to have thought of ethical facts as what Darwall (2003) calls “agent-

character; that is, the fact that they seem to be saying of some agent that he ought to do something”(1968: 102).

²⁷ I say no *particular* ontological commitments because practical judgments and claims might presuppose general ontological commitments to things like the existence of agents.

neutral” facts about what has intrinsic value. Moore claims that what is intrinsically valuable is what “ought to exist for its own sake” and thus judging that one ought to perform a particular action is, for Moore, to judge that performing this action will produce more overall goodness: “To assert that a certain line of conduct is, at a given time, absolutely right or obligatory is obviously to assert that more good or less evil will exist in the world, if it be adopted than if anything else be done instead”(1903: 25). So, for example, when one judges that one ought to be charitable one is judging that actions of charity produce something of intrinsic value, meaning something that is good independent of all of its extrinsic properties, including things like whether it is pleasant, culturally encouraged, or a goal set by the agent who ought to be charitable.

So maybe the motivation for the nonnaturalists’ hypostasizing the objects of ethical knowledge is that this makes ethical claims, so to speak, answerable to an ethical reality that is objective and independent of human practices and reasoning patterns. In short, positing ethical facts is a way to secure the conclusion that ethical dictates are universal. In this sense, Moore’s view of ethical claims is isomorphic to a realist view of prosaically natural claims: as the latter are answerable to an independent objective natural reality, the former are answerable to an independent objective ethical reality. However, I think the way Regan defends nonnaturalism against nonreductive a posteriori naturalism belies this attempt to underwrite universalism.²⁸ For according to him the reason we should posit nonnatural ethical facts in the first place is that ethical judgments are distinctively practical in that we make ethical judgments only in the course of practical deliberation and ethical knowledge arises out of practical rather than speculative deliberation. However, there is an important disanalogy between practical and

²⁸ Although it’s worth noting that it doesn’t belie all attempts. One can be a universalist without positing objective ethical facts of any kind. One just needs to think there are universal principles governing any piece of practical reasoning. Even if practical reasoning is agent-relative, it might be relative to each agent in a universal way.

speculative deliberation having to do with the fact that the goal of practical reasoning is (some kind of) action while the goal of speculative reasoning is (some kind of) belief.

To see this, notice that both practical and speculative reasoning have what Velleman (1996/2000) calls “formal aims”, i.e. aims defined completely in terms of, or in terms that essentially depend on, the very concept of being the object of the relevant activity. Velleman gives the example of winning as the formal aim of competitive games, because winning just is what it is to succeed at a competitive game. Similarly, the formal aim of practical reasoning is to perform the action for which one has best reasons, while the formal aim of speculative reasoning is to believe the proposition that one has best reasons to believe. This, of course, doesn’t specify what constitutes an ‘action for which one has best reasons’ or ‘belief for which one has best reasons’ any more than the claim that winning is the aim of competitive games specifies what constituted ‘winning’—this is why these are the formal aims. However, in addition to its formal aim, speculative reasoning seems to have a substantive aim which can be derived from the concept of being the goal of speculative reasoning. The goal of speculative reasoning is (some kind of) belief, and belief is a propositional attitude that is—as it is sometimes put—“aimed at the truth”. This is a somewhat vexed metaphor, but I take it to come to at least the fact that beliefs are correct iff true, where what it is for a belief to be true is necessarily agent-neutral, i.e. something that doesn’t depend on the attitudes or endorsements of any given agent.

If Regan wants to win the objectivity of ethical facts by positing an independent ethical reality that we only have ‘access’ to through the operation of practical reasoning, it would seem that he needs there to be a substantive aim of practical reasoning that is analogous to the way truth is the aim of theoretical reasoning. But this won’t work. There is much debate about whether there is a unified correlative substantive aim of practical reasoning; however, this much seems clear: there is no substantive agent-neutral aim that can be derived from the concept of being the goal of practical reasoning. For, unlike the concept of ‘belief’, the concept of ‘action’

doesn't carry a non-formal standard of correctness with it. Regan might suggest that action is aimed at intrinsic value just as belief is aimed at truth; in neither case do we have anything more than our always-evolving complete system of, on the one hand, considered critically resistant beliefs, and on the other hand, considered critically resistant desires to work from, but in both cases we are aiming at something that is agent-neutral. However, I think we can turn Regan's own notion of the "desire-transcending question" and something like Moore's open-question argument against him here. As Darwall (2003) has argued, it is easy to close (in a Moorean sense) the question "p is X but should I believe p?" by substituting 'true' in for X to thereby get an agent-neutral standard for belief; but an equivalent substitution in the question "a is X but should I do a?" with 'intrinsically valuable' does not generate a similarly closed question. For example, it may be intrinsically valuable for me not to return your gun, but, given that I promised to return it, it is certainly not conceptually confused to wonder whether I should do so. Darwall puts the point in terms of freedom:

As agents, we have a freedom that we simply do not have as believers. We can step back from any perspective given us even by our most critically informed desires and sensibly ask whether what we have reason to do is just to promote valuable states (as they will seem from our desires' perspective). And we can assert without contradiction or conceptual confusion that this is not so, that there are other, agent-relative reasons. It follows that unlike theoretical reason, practical reason is responsible to no external goal or standard which logically closes sensible deliberation. Reasons for acting can be grounded nowhere but within norms of free practical reflection itself.(Ibid.: 485)

I think Darwall overstates the conclusion here. Some reasons for acting might be grounded in something like agent-neutral intrinsic value; however, I think he is right that no external standard logically closes sensible deliberation and at least some reasons for action seem to be grounded in the agent-relative strictures of practical reasoning itself. In my opinion, this is a reflection of the mundane fact that in engaging in speculative reasoning, one attempts to reach a correct conclusion about what to believe, and if correct this conclusion is correct for anyone. However, in engaging in practical reasoning, one attempts to reach a correct conclusion about

what to do, and even if one's conclusion is correct this conclusion is not necessarily correct for anyone.

We can see the problem for Regan here by contrasting ethical facts with other sorts of fact that are sometimes thought to be nonnatural. Facts about abstract objects such as numbers and sets can seem to be incapable of investigation by the natural sciences; nonetheless, we want to say that mathematical facts are universal and that we come to know them in a way noticeably different from the way that we come to know prosaically natural facts. Because of this, ethical nonnaturalists often try to bolster their position by pointing to mathematical facts as a partner in the putative crime against the doctrine of naturalism. Such partners-in-crime arguments are only as strong as the intuition that the putative partner is truly innocent of the putative crime; and so one who sees mathematical facts as just as or more problematic than ethical facts will not be swayed by the nonnaturalist's appeal to mathematical facts as a partner in crime. However, even if we allow that mathematical facts are completely unproblematic, notice that ethical facts, conceived of as known through the exercise of practical deliberation, exhibit the following disanalogy with mathematical facts: they are action-guiding

Now, I don't think it would do any good for the nonnaturalist to react to this problem by jumping ship and endorsing a posteriori naturalism. I think nonnaturalists are right that ethical judgments and knowledge play a distinctive practical role in our cognitive economy—this is force of the thought that ethics is peculiarly action-guiding—and this implies that ethical knowledge is achieved in a way that is discontinuous from the way knowledge of prosaically natural facts is achieved. So, in my view, the nonnaturalist's epistemological lesson is correct. It's just that this epistemological lesson undermines rather than underwrites the ontological conclusion that nonnaturalists want to draw from it. That is, the discontinuities between ethical and natural knowledge show that there's no reason to admit objective and independent ethical

facts into our ontology. But if that's the case, then we cannot understand the meaning of ethical claims in terms of their being descriptions of an objective universal realm of ethical facts.

So, in the end, even if Regan's transcendental argument for the claim that ethical knowledge emerges out of our conception of ourselves as practical reasoners, i.e. agents, provides good reason for thinking that ethical knowledge is *sui generis*, treating this as knowledge of a *sui generis* realm of nonnatural facts seems to be unnecessary for capturing the *sui generis* character of ethical knowledge. In the following section, I shall consider a different way that some metaethicists have sought to respond to Mackie. This is to give up the assumption that ethical claims purport to be universally true.

III. RELATIVISM

So far, we have explored universalist ways of working out the idea that the meaning of ethical claims is to be given in terms of what fact they state or describe. I think there is a *prima facie* advantage of these accounts: they capture the intuition many philosophers have had that ethics is fundamentally universal. Ethics seems to be more than a matter of local cultural norms or personal preference. However, these universalist accounts were problematic in various ways. So this (combined with the prominence of relativist views throughout philosophy and in popular culture) encourages the exploration of relativist attempts to account for the meaning of ethical claims. Specifically, I want to consider versions of what is usually called *metaethical relativism*.

Metaethical relativists hold, in general, that the meaning of ethical claims is implicitly relative in that even though there is of the ascription of a n -place predicate to a subject, their deep logical form is actually of the ascription of a $n+1$ place predicate. For example,

(6) Joe ought to care for his children,

may seem to mean that Joe's caring for his children has the property of being morally obligated; however, a metaethical relativist will maintain that there is actually a further variable argument place in moral predicates that is semantically tied to some independent source of moral standards. So, for instance, a cultural relativist might argue that (6) means not that Joe's caring for his children is morally obligated, *simpliciter*, but that Joe's caring for his children is morally obligated *by Joe's culture's moral norms*.²⁹

Why hold such a view? There are perhaps several reasons, but one is obviously that it provides an answer to the challenge posed by Mackie's argument from queerness. Mackie's argument turns (in part) on the view that the only respectable facts that one can appeal to in an account of reality are natural facts, where a fact is a natural fact (roughly) if and only if it can be the object of empirical investigation and confirmation through the methodology of wide-equilibrium constitutive of science. Construed as facts about universal values, however, it seems initially difficult to see how ethical facts could be natural facts. A relativist will argue that ethical facts can be easily seen as natural facts if they are recognized to be relative facts. For example, the personal preferences relativist holds that the fact that makes my claim, "War is immoral" true is a fact about my preferences—perhaps, that I dislike war in a particularly moralistic way. And he will argue that as long as facts about personal preferences are natural, there is no problem with squaring ethical facts with naturalism. Thus, a commitment to naturalism can support relativism.

A second reason often cited for being a relativist is that relativism can best explain what we observe when we take up a, so to speak, anthropological stance towards moral discourse. More so than in most areas of discourse, there seems to be widespread difference of opinion. For many ethical sentences S, we can find people who will sincerely assert S and people who will

²⁹ A metaethical cultural relativist of a different sort might argue that ethical claims are relative to the moral standards of the claimant's culture. Below I'll discuss the difference between what might be called speaker-relativism and agent-relativism.

sincerely deny S. This is, of course, true in other areas of discourse, but there it seems to be a result of some lack of information or cognitive defect on one or the other side; and while differences in opinion about ethical matters is surely also sometimes a result of some lack of information or cognitive defect, there also seem to be recalcitrant differences of opinion that survive correction for lack of information or cognitive defect.

As I mentioned in the introduction above (page 7), one way to explain such recalcitrant differences in opinion is to posit an implicit relativity in the logical form of the different claims. In this way, when the value of this deictic element is fixed different for the two claims, they can be seen as expressing different but not incompatible propositions. Now, there is a lot room here for different kinds of relativism and thus a lot of work for a relativist to do to defend her specific choice of deictic element: is it something subjective or something intersubjective? is it determined by features of the speaker or features of the evaluated agent? Whichever way a relativist goes here, however, she can claim some support for her view due to its ability to explain why differences in ethical opinion seem to be particularly recalcitrant.

But how exactly should the relativist spell out her position? There are two salient axes along which different families of relativism could be differentiated. First, ethical claims could be relative to things like cultural norms, legal codes, personal values, brute preferences; second, the relevant one of these for determining the meaning of an ethical claim could be linked to the speaker of the claim or to the implicit or explicit moral agent the claim is about. The above example about (6) was an example of cultural agent-relativism because the meaning of the claim was thought to be given in reference to the moral norms of the evaluated *agent's culture*. However, other versions of metaethical relativism can be generated by changing these two variables. For example, a cultural *speaker*-relativist would say that this claim means that Joe's caring for his children is morally obligated by the moral norms of my [the speaker of the claim] culture. (Of course, these might come to the same thing if the speaker's culture is the same as

Joe's culture.) And personal values agent-relativist would say that this claim means that Joe's caring for his children is morally obligated by his [Joe's] own personal values.³⁰

Most seriously defended versions of relativism are versions of speaker-relativism. In fact, the only exception I am aware of is Harman (1975), who argues that the core use of the 'moral ought' is to make what he calls "inner" judgments, which he thinks are relative to both a feature of the *agent's* motivational profile and an endorsement of the *speaker*.³¹ On this view, the content of (6) is (roughly) "Joe's taking care of his children is obligated by his [Joe's] reasons for acting." This is not to say that one who is committed to (6) is committed to Joe's being actually motivated to take care of his children—for he may not be actually motivated to do what his reasons obligate him to do—but it is to say that (6) commits one to the view that Joe has (justifying) reason to take care of his children.³² About such "inner" judgments, Harman writes, "[they] imply that the [evaluated] agent has reasons to do something"³³(Ibid.). This is straightforwardly a commitment the doctrine I've been calling rationalism. However, Harman takes this rationalism to support a complex form of relativism because he thinks that what it is

³⁰ As we'll see in chapter 3, there are metaepistemological relativist views that parallel various versions of metaethical relativism; and one of the most prominent of these is a version of *agent*-relativism rather than *speaker*-relativism.

³¹ He thinks that there are also "external" judgments that do not imply anything about the evaluated agent's motivational profile; and he uses this distinction to explain why it can seem correct to say, e.g., that the Vikings were savage, but incorrect to say, e.g., that they morally ought not have raped and pillaged. The first is an external judgment that implies nothing about the Vikings motivational profiles, while the latter is an inner judgment that (implausibly) implies that the Vikings could have been motivated by the pain and suffering they caused to cease raping and pillaging. Here I am concerned primarily with such ought-claims, so I'm happy to concede the distinction for the sake of argument.

³² In reading Harman as committed to rationalism, I've read his claim about having reasons as being about *justifying* reasons. Harman is not completely clear on the difference between justifying and motivating reasons; however, if these claims are read as being about motivating reasons, then Harman is committed to the highly implausible view that I above called "motivation agent internalism". So generosity supports my reading.

³³ In addition, they imply that the speaker endorses these reasons. This may make Harman's view doubly relativist in that the content of ethical claims is relative to some aspect of the agent and to some aspect of the speaker. In what follows, for the sake of contrast and simplicity, I will suppress the sense in which Harman might also be a speaker-relativist and focus on the argument he offers from rationalism to agent-relativism.

to think that one has reason to ϕ is to endorse considerations that can actually come to motivate the agent to ϕ ; and he thinks that these considerations vary depending both on who is making the ethical claim and who the claim is a claim about. This is because he thinks that “reasons for action must have their source in goals, desires, or intentions [of the agent]”(Ibid.: 9). And he thinks that making an inner judgment that ethically evaluates someone presupposes that they implicitly intend to follow a common code of action.

There is clearly a sense in which this form of metaethical relativism brings resources for an explanation of the action-guiding character to ethical judgments that universalist naturalists do not have. So, a third more complex motivation for relativism is that it can help to explain the action-guiding character of ethical judgments. However, as we saw above, the action-guiding character of ethical claims seems to come (at least) to the fact that two different thoughts are intuitive and highly plausible. First was the rationalist idea that if it is true that ϕ -ing in C is good, then one has a justifying reason to ϕ when in C. Second was the (weak) internalist that, insofar as one is normal, one who judges that ϕ -ing in C is good will be ipso facto fallibly disposed to be motivated to ϕ when in C. As far as I can tell, Harman’s relativism only captures the first.

He does capture a connection between ethical judgments and motivations to action, but it’s the evaluated agent rather than the speaker. I think this is because he implicitly accepts a third doctrine related to action-guiding that needs to be distinguished from rationalism and internalism. This is the (confusingly labeled) doctrine called *reasons internalism*, which is the view that in order for one to have a (justifying) reason to ϕ in C, one must either already have a motivating reason to ϕ when in C or there must be a rational path by which one can come to have a motivating reason to ϕ when in C. Reasons internalism doesn’t have anything in particular to do with ethical claims and reasons, but is rather a general thesis about a necessary condition on justifying reasons. We can distinguish between weak and strong versions of

reasons internalism depending on the interpretation of “rational path”. Some, such as Korsgaard (1986), endorse a weak version in that they think that if there really is a (justifying) reason to ϕ in C, than *any* rational agent can come to be motivated to ϕ when in C. That is, “rational path” is interpreted as being in principle open to any rational agent. Others, such as Williams (1980/1981), endorse a stronger version in that they think that (justifying) reasons are essentially nonuniversal. That is, “rational path” is interpreted as depending essentially on the, so to speak, motivational starting point of the agent.

I suspect that Harman endorses rationalism and implicitly accepts a strong version of reasons internalism, which is what leads to his relativist view. But notice that if one rejected the strong version of reasons internalism, then much of the motivation for Harman’s relativism would be removed. In my opinion, the debate over reasons internalism is difficult and vexed, and this is precisely because we don’t have strong intuitions to work from here but only developing normative theories. As a result, I think Harman’s relativism rests on a pretty contentious issue, and a more plausible form of relativism would remain neutral on the issue of reasons internalism.

It’s easy to see how a better form of relativism can be developed in order to avoid the issue of reasons internalism and to explain internalism. For even as weak as we’ve made internalism, it’s striking that there’s a connection between a form of judgment and motivation at all. Dreier puts the point well:

...how could even modest internalism possibly be true? It is puzzling how there could be beliefs, the mere having of which would be logically connected to a motivation. It is not that our folk psychology conceives of beliefs and desires as completely independent, there is no mystery how a belief that by doing A I will do B could lead to my motivation to do A—provided that I am motivated to do B. The worrying feature of internalism is that ...it says that there are beliefs the mere having of which normally establishes the presence of a motivation without depending on any other prior motivation.(1990: 14)

Dreier argues that the best explanation of the truth of internalism will build a connection to motivation into the truth-conditions of the ethical claims. There are many ways one might do

this, depending on what one takes to motivate ethical action, but in each case the basic idea is to make the truth conditions relative to the appropriate motivational feature.

It's important to notice that this argument for relativism from internalism supports only one broad family of relativism—it only supports forms of speaker-relativism, which probably explains the popularity of this position in contrast to agent-relativism or some mixed view like Harman's. Westermarck (1932), for instance, defends the view that ethical statements are relative to the speaker's dispositions to manifest a certain attitude in certain situations. So, he would say that the fact stated by (6) is that Joe's not caring for his children belongs to the class of actions towards which the speaker of (6) tend to take up an impartial attitude of angry resentment. If such attitudes motivate the appropriate disposition to action, then a form of internalism falls out of this account of the content of ethical claims.³⁴

Some have worried, however, that simple preferences and the attitude of angry resentment aren't rich enough to underwrite a recognizably moral code. Thus, philosophers such as Wong (1984) and Dreier (1990) defend a form of metaethical speaker-relativism according to which the implicit variable in ethical claims is fixed by *all* of the speaker's affective and motivational states (somehow) idealized into what they call a "practical system". Thus, they would say that the fact stated by (6) is that Joe's caring for his children is highly valued by standards of the speaker's practical system. These are subjectivist versions of metaethical speaker-relativism because they view the implicit variable in ethical claims as fixed by some subjective feature of the person making the claim.

³⁴ One minor complication: Even if the truth-conditions of ethical claims are relative to the speaker's motivational states, speakers could make ethical claims without having the states requisite for motivation—for they could be insincere or simply mistaken about what motivational states they have. This can be finessed by construing the internal connection between ethical claims and motivation as follows: one who sincerely makes an ethical claim must take himself to be somewhat motivated to act in its accord. Compare Ridge (forthcoming-b).

Is this argument from internalism to subjectivist speaker-relativism good? If internalism is correct and what it means to be normal is (at least in part) that one is motivated by one's desires or the values of one's practical system, then subjectivist speaker-relativism can underwrite and explain why this is so. But what about rationalism? If we relativized claims of the form " ϕ -ing in C is good" to the speaker's desires or other subjective states, one might reasonably wonder whether speaker-relativism wins an explanation of internalism at the cost of foreclosing any plausible explanation of rationalism. For, although the fact that I like ϕ -ing in C *can* give me a reason to ϕ in C, facts of this form cannot be intrinsically reasoning-giving. For one thing, they clearly don't have to give other people who dislike ϕ -ing in C a reason to ϕ in C. And for another, when we recall that we're talking about justifying reasons here, it seems fairly clear that this fact doesn't have to give me a reason—as is the case when my preferences are immoral.

Many philosophers will say that this is a reflection of the objective character of ethics. When one of my ethical claims is true, the reason-giving potency of this truth goes not just for me but for others too. However, it's important to notice that we can endorse this idea without giving up on speaker-relativism. Prinz (forthcoming), for instance, defends an inter-subjectivist version of speaker-relativism, according to which the content of one's ethical claims is implicitly relative to the emotional reactions of a normal member of one's community under circumstances of full information. So, for him, (6) would be thought to state the fact that Joe's caring for his children evokes emotions in the positive range among normal members of the speaker's community under circumstances of full information. This view can explain why rationalism is true as long as one takes the positive and negative emotional reactions of members of one's community under circumstances of full information to track reasons for and against an action. We may hope for an even more objective conception of reasons for action; but Prinz will insist that this hope is in vain—it leads to bad metaphysics and it is just as good to explain reasons for action in terms of the sensibilities of fairly idealized observers. The idea is

that one has reason to have the motivation appropriate to one's ethical claim because normal members of one's community under circumstances of full information would have this motivation.³⁵ This is a sort of deference in moral sentiment: it's always possible that my own emotional reactions to a particular situation are skewed because I am not fully informed or not normal, and this possibility of error should be reflected in the content of my moral claims; but that doesn't require endorsing premise (1) of Mackie's argument from queerness, since we can capture it with an inter-subjectivist form of speaker relativism.

But we might fear that Prinz wins his explanation of rationalism at the cost of foreclosing any plausible explanation of internalism. For although his view connects the content of ethical claims, so to speak, to motivationally potent psychological states, the psychological states he thinks relevant to the content of the claim—viz. those that would be had by normal members of one's community *under circumstances of full information*—are not necessarily states had by the claimant. But here I think Prinz can get much closer to capturing internalism than subjectivists get to capturing rationalism. He can argue that internalism doesn't require that one who makes an ethical judgment *necessarily* have any particular motivational state. It only requires that insofar as one is normal one be fallibly disposed to be appropriately motivated. And part of being fallibly disposed might just be being more likely to be appropriately motivated as one gains more and more (relevant) information. Thus the motivational state relevant to Prinz's characterization of the meaning of ethical claims is a sort of idealization out of the dispositions of those normal people who make the claims.

³⁵ This is not, by the way, the way Prinz himself argues for his version of relativism, although I think he would accept this as a supporting argument for his metaethical view. Without getting into the details here, let me just mention that he adopts an information-theoretical semantics (*a la* Dretske (1981)) and argues that moral concepts have the natural function of tracking the property of evoking certain (ethical) sorts of emotional reaction among ideally informed members of a community. This of course opens him up to objections about informational-theoretical semantics, which is why I don't treat this as one of the possible general arguments for relativism in the text above.

I'm not sure that this completely does the trick, but I think Prinz's inter-subjectivist form of metaethical speaker-relativism can do a better job of capturing the action-guiding character of ethical judgments within a descriptivist framework than the subjectivist versions.

So, three key motivations for metaethical speaker-relativism emerge from these reflections. First, the view offers a way out of Mackie's argument from queerness, so insofar as its conclusion is nonnegotiablely unacceptable, this tells in favor of relativism. Second, the view can explain the existence of recalcitrant differences of ethical opinion. Third, some versions of the view go most of the way towards explaining the action-guiding character of ethical judgments. I don't think these are knock-down arguments for metaethical speaker-relativism, because they could also support other positions, but I do think they each mobilize real and powerful considerations that any account of the meaning of ethical claims should take seriously.

Despite these advantages, there are well-known problems for metaethical speaker-relativism stemming from the fact that it violates a number of ordinary semantic intuitions. For example, notice that speaker-relativists have to accept as legitimate pieces of ethical discourse that would strike any ordinary speaker as illegitimate. To see this, consider how Westermarck would analyze the following claim,

(7) Taking the money is wrong; so don't do it.

He would say that the first part means,

(7') Taking the money belongs to the class of actions towards which I tend to take up an impartial attitude of angry resentment.

But now, consider how the analysis goes if someone replied to the original claim with,

(8) Taking the money is not wrong, so I'm going to do it.

Westermarck would say that the first part of this reply means,

- (8') Taking the money does not belong to the class of actions towards which I tend to take up an impartial attitude of angry resentment.

However, the referent of the first person pronoun 'I' in (7') and (8') are different, which means that the original speaker could cogently—if Westermarck is right about the meaning of these claims—continue the dialogue:

- (9) Ok, but taking the money is wrong, so don't take it.

Now, from the perspective of ordinary semantic intuitions, this rejoinder sounds like a very odd thing to say in a conversation after the (7) and (8) have been uttered. Moreover, (7) seems to present a valid inference (whose nature, of course, needs to be explained), but replacing the antecedent of (7) with (7') seems to turn it into an invalid inference. Opponents of ethical relativism suggest that these odd results reveal that the meaning of ethical claims cannot really be relativized to features of the speaker at all. If relativism is right and we manifest at least vague awareness of this semantic fact, we shouldn't find (9) so odd or think the validity of (8) is ruined by replacing the antecedent with (7'). Notice how a similar exchange, where a relativistic semantics is uncontroversial, doesn't exhibit this oddness. On the telephone in different cities, two people might exchange the following sentences:

- (10) It's warm here, so I won't bring a jacket unless you think I'll need it.

- (11) It's cold here, so I think you should bring a jacket when you come.

- (12) Ok, but since it's warm here, I'll pack my jacket in my checked bag.

That is, a critic will insist that our semantic intuitions about the oddity of (9) in contrast to our semantic intuitions about the appropriateness of (12) reveal that something is wrong with the relativistic analysis of (6) and (7). Moreover, if we replace the antecedent of (11) with

- (10') It's warm in location l,

that doesn't intuitively turn a valid inference into an invalid one.

So even after long discussion and debate, well-informed and apparently rational people can continue to express different ethical opinions; the relativist claims that this diversity does not represent genuine disagreement because apparent disputants actually mean different things by their ethical claims. Opponents will suggest even if this explains the intransigence of the difference of opinion, it does so in a way that then violates the speakers' own semantic intuitions about the nature of their claims; specifically they will take themselves to be disagreeing, and the relativist must reject this semantic self-interpretation. (Moore 1912) It may be true that in some cases of apparent disagreement, reflection will lead us to the view that there is no real disagreement. However, for most real cases of differences of ethical opinion, most have a strong intuition that the *difference* is a genuine *disagreement*.

Why do these countervailing semantic intuitions represent a problem for the metaethical relativist? Because she provides an account of the meaning of ethical claims that almost all competent users of these claims would, upon reflection, reject. At its core, I think this objection is a complex application of the Moorean strategy of mobilizing semantic intuitions against proposed analyses of ethical claims. The speaker-relativist claims that an ethical claim of the form

(13) A ought to X

means, roughly,

(14) A's doing X is validated by N,

where N is a variable for which a moral standard determined by some nonuniversal feature of the speaker can be substituted. But most competent users of ethical claims would, upon reflection, deny the semantic equivalence of instances of (13) and (14).

As we've already seen in the discussion of a posteriori naturalism above, a common response to this Moorean strategy is to argue that these semantic intuitions aren't probative with respect

to the correct analysis of ethical claims. This is because the fact described an instance of (13) may be identical to the fact described by the corresponding instance of (14) even though this identity is not borne out in competent speaker's intuitions about the semantic value of their claims. Thus, a relativist might try to respond to the line of objection we've been considering by claiming that his account is a metaphysical account of the facts described by ethical claims rather than an account of the meaning of ethical claims. It is in this vein, I think, that Harman writes,

Einstein's Theory of Relativity does not involve a claim about meaning or about what people intend to be claiming when they make judgments about an object's mass. The point is, rather, that the only truth there in this area is relative truth. Similarly, the moral relativism I will argue for is not a claim about what people *mean* by their moral judgments. Moral relativism does not claim that people intend their moral judgments to be "elliptical" in the suggested way; just as relativism about mass does not does not claim that people intend their judgments about mass to make implicit reference to a spatio-temporal framework.(Harman and Thomson 1996: 4)³⁶

I don't want to evaluate whether the resulting metaphysical thesis is correct, because all that is important for my purposes is to notice two fishy things about the move on the relativist's part. First, it effectively abandons the attempt to respond to Mackie's argument from queerness; in fact, the resulting position is *very* similar to Mackie's view that ordinary positive ethical claims are false, but that there's a closely related relativized claim that we could make instead that would be true.(Mackie 1977: ch. 2) Second, Einstein's theory of relativity is clearly not meant as a semantic theory; however, metaethical relativism should be, to put it in Harman's own words, "a soberly logical thesis—a thesis about logical form"(1975: 3). But if this is not then a thesis about the content or meaning of ethical claims, then it gives up on the descriptivist approach to accounting for the meaning of ethical claims; for a relativist who makes this move can no longer claim to be giving an account of the meaning of ethical claims in terms of the facts that they

³⁶ Although, as I discussed above, Harman's view appears to be a version of agent relativism, obviously a speaker relativist could attempt the same move in response to the semantic objection I have been discussing.

state. Thus, insofar as one wants to give a descriptivist account of ethical claims, this relativist move in response to the objections is unavailable.

Another response to the line of objection we've been considering, is to argue that our semantic intuitions about (7)-(9) and (13)-(14) are simply not to be trusted. Intuitions of genuine disagreement might be explained away by claiming that the value of the implicit variable to which ethical claims are relative is so often something shared by both speakers that the implicit variable is very deeply hidden from ordinary semantic intuitions. Similarly, the intuition of semantic difference between ethical claims and their relativistic analyses might be explained away by pointing out that the most plausible analysis of almost all ordinary claims posits implicit variables that do not necessarily manifest themselves in ordinary semantic intuitions. For instance, the most plausible analyses of attributions of flatness construe them as implicitly relative to variable standards determined by something like the implicit purposes for which the evaluated surface is to be used. But such semantic relativity doesn't necessarily show up if you ask ordinary speakers whether sentences of the form,

(15) X is flat,

mean

(16) X is level enough for present purposes.

Why think that ordinary intuitions about the synonymy of ethical claims and their relativistic analyses are any more probative than ordinary intuitions about the synonymy of instances of (15) and (16)?

This is a natural partners-in-crime response to the objections, but it invites an important rejoinder. In constructing an account of the meaning of ethical claims, one can always deny ordinary intuitions that don't comport with the account, but this denial comes at the cost of requiring a plausible explanation of why ordinary speakers have the putatively mistaken

intuitions. Of course, how big this cost is will depend on how strong the intuitions are that the account violates. I think the intuitions violated by ethical speaker-relativism are much stronger than those violated by in plausible cases of semantic relativity such as flatness attributions.

Even if ordinary semantic intuitions don't explicitly manifest awareness of the implicit relativity in flatness attributions, ordinary speakers can quickly be brought to see the relativity; and this is not true in the ethical case. To see this, find two thoughtful and patient ordinary speakers who are familiar with a number of the same surfaces, tell the first that you need a place to set up your computer, and then ask her whether various surfaces (that she is familiar with but she knows you are not familiar with) are flat. Once she says that some surface is flat, ask the other person whether the indicated surface is flat enough to do a very sensitive scientific experiment for which a flat surface is absolutely crucial. When he says no, point out that he seems to have contradicted the first person's earlier claim. More often than not, I submit, patient ordinary speakers will respond to such an apparent contradiction by appealing to the purpose relativity of flatness attributions. The situation seems clearly disanalogous when it comes to ethical claims. If one person says "X is wrong", another person—even another person from a different culture—says "X is not wrong", and they are made aware of this apparent contradiction, I doubt that they will appeal to implicit relativity of their claims in response. Much more likely is that they will try to give reasons for their view and against the apparent disputant's view.

The point here is not whether these differing responses would be correct. Rather, since I am now concerned to show only that the relativist's partners-in-crime response to the two objections we've been discussing is not plausible, it's enough to notice that these typical responses would be different, which I take to indicate that the crime against ordinary semantic intuitions committed by the metaethical relativist is much graver than the crime committed by similar relativistic accounts of flatness attributions.

IV. THREE CHALLENGES

In this chapter, we have explored the three most prominent families of descriptivist accounts of the meaning of ethical claims—naturalism, nonnaturalism, and relativism. There were positive motivations for each view, but each also suffered serious drawbacks. I think this discussion reveals three complex challenges to developing an adequate account of the meaning of ethical claims

First, there is what I shall call *Mackie's challenge*. Any account of the meaning of ethical claims needs to show how to respond to Mackie's argument. Mackie argues that putative ethical facts are peculiarly action-guiding, which I cashed out in terms of a commitment to rationalism and internalism. Strictly speaking, these are theses about ethical judgments, but in the context of descriptivist theories of the meaning of ethical judgments, they can be transmuted into theses about the intrinsic reason-giving and magnetizing character of putative ethical facts. Mackie insists that such facts would be too queer to admit into our ontology. Naturalists, nonnaturalists, and relativist can all be seen as offering different ways to respond to Mackie's challenge. Naturalists deny premise (2) of the argument from queerness—that is they deny that ethical facts really do have the action-guiding character that makes them seem queer to Mackie. Nonnaturalists deny premise (4) of the argument from queerness—that is they deny that nonnatural facts with an action-guiding character really are too queer to be admitted into our ontology. And relativists deny premise (1) of the argument from queerness—that is they deny that ethical claims purport to be objectively true in that their truth is independent of any policies, preferences, or choices of particular groups of people. As we've seen, however, each of

these three strategies for responding to Mackie incur costs, and these costs generate the further two challenges.

Second, there is what I shall call the *action-guiding* challenge. I suggested that naturalists' denial of the intrinsic action-guiding character of ethical judgments turns on weak arguments that miss the point and/or attack straw men. The rationalist idea was that if it's true that ϕ -ing in C is good, then one in C has (justifying) reason to ϕ ; and the internalist idea was that insofar as one is normal when one judges that ϕ -ing in C is good one is ipso facto fallibly disposed to be motivated to ϕ in C. These two ideas are intuitive and highly plausible. Any account of the meaning of ethical claims needs to respect them and preferably provide resources for explaining why they are true. Nonreductive a posteriori naturalists cannot do this, which is a serious drawback to the account. Nonnaturalists purport to be able to do this, by in essence pulling ethical facts out of the natural realm and making them nonnatural. However, it is mysterious how that really helps; for as far as meeting the action-guiding challenge goes, saying that ethical facts are nonnatural is by itself about as helpful as saying that ethical facts are special in just the right way. Moreover, being special in just the right way is precisely what generates Mackie's challenge, so structurally nonnaturalists seem to be in a bind. Shafer-Landau and Regan seek to quell Mackie's worries by generating moral epistemologies that are distinct from standard empirical epistemologies, and I think Regan's version of this makes moves in the direction of answering the action-guiding challenge; however, as we saw, objective ethical facts turn into a fifth wheel on his account, and so it is unclear that his nonnaturalism offers a genuine version of descriptivism. Furthermore, although some relativist accounts offer a very good account of the truth of internalism, they have a hard time capturing the rationalist idea unless they construe ethical claims as implicitly relative to something with justificatory force. If this were the only problem with relativism, I'd be inclined to think that relativism beats out naturalism and nonnaturalism vis-à-vis the action-guiding challenge, and that we should reexamine the validity

of the rationalist idea that a commitment to an ethical truth commits us to reasons for acting in a particular way. However, the discussion of relativism revealed a third challenge.

Third, there is what I shall call the *relativist's challenge*. In our discussion, a number of related considerations against relativism emerged once we considered how plausible it is as a semantic view about the content of ordinary ethical claims. These problems all turn on the fact that the relativist is forced to violate ordinary semantic intuitions about the genuineness of ethical disagreements; and this violation is theoretically more costly than similar violations by relativistic analyses of other sorts of claims. I think it would be wrong to treat this as irrefutable evidence against metaethical relativism. After all, these drawbacks can be seen as the flipside of an obvious advantage: Relativism can explain the presence of particularly recalcitrant differences in ethical opinion. This means that we want an account that can explain the presence of particularly recalcitrant differences in ethical opinion in a way that doesn't treat the difference as mere difference and thereby disrespect the intuition that it represents disagreement.

2 EXPRESSIVIST ACCOUNTS OF ETHICAL CLAIMS

In the previous chapters I explored the three most prominent general attempts to explain the meaning of ethical claims in terms of what facts they state or purport to describe. We saw that each of these is motivated by some important considerations; however, each of them also suffer some serious objections. I used this discussion to generate three challenges on an acceptable account of the meaning of ethical claims: Mackie's challenge, the action-guiding challenge, and the relativist's challenge. The fact that none of the three prominent descriptivist accounts satisfactorily meet all three challenges should, I think, encourage consideration of some alternatives. So, in this chapter, I begin my exploration of nondescriptivist alternatives to these descriptivist accounts.

The descriptivist accounts explored in the previous chapter shared the common organizing principle that the meaning of ethical claims is to be explained in terms of what facts they state or purport to describe. A prominent way to develop nondescriptivist alternatives is to replace this organizing principle with the principle that the meaning of ethical claims is to be explained in terms of what they express. This is why the most prominent nondescriptivist accounts of ethical claims are typically called "expressivist". But we should be careful about this label, for there are two senses in which even a descriptivist can accept that the meaning of ethical claims is to be explained in terms of what they express. First, a descriptivist can say that an account of what

fact is stated by ethical claims *just is* an account of the proposition “expressed” by the claims. Second, a descriptivist can say that ethical claims “express” beliefs, whose content is that a particular fact obtains, and then go ahead and transpose their account of ethical facts into an account of the type of beliefs expressed by ethical claims.

It is for this reason that expressivists typically insist that ethical claims express something of an importantly different kind from descriptive claims. One reason some philosophers have found the expressivist approach attractive is that it provides a response to Mackie’s challenge that does not require vindicating the existence of ethical facts. In effect it is a rejection of Mackie’s premise (3) in the argument articulated on page 14 above; that is, expressivists deny that one who makes an ethical claim is thereby committed to the view that there exists a corresponding fact. This allows expressivists to (a) affirm the legitimacy of ethical discourse, (b) deny the existence of nonnatural facts with a *sui generis* action-guiding character, and (c) deny that ethical claims state either relative or nonrelative natural facts. Their characteristic move to this end is to argue that ethical claims express something that, so to speak, lacks descriptive aspirations. Usually, therefore, different expressivist accounts are separated according to the specifics of their accounts of what ethical claims express and why, whatever this is, lacks descriptive aspirations. For instance, emotivists such as Ayer (1946) and Stevenson (1937) argue that ethical claims express a sort of emotion, while quasi-realists such as Blackburn (1994, 1998) and Gibbard (2003) argue that ethical claims express a sort of practical attitude.

Although this way of distinguishing the various extant expressivist positions is correct, I think it obscures an importantly prior issue—viz. the sense of ‘express’ deployed in each of these claims. So, here, in this chapter I want to consider this matter in some detail. Specifically, I shall argue that traditional emotivist versions of expressivism deployed two different senses of ‘express’, which helps them to meet a number of the desiderata outlined at the end of the previous chapter. However, I think this also opens them up to a number of serious objections

that contemporary quasi-realist versions of expressivism avoid by deploying only the second of these senses of 'express'. Unfortunately, as we shall see, this strategy both makes it hard for them to meet some of the challenges outlined at the end of the previous chapter and opens them up to an objection with their way of responding to Mackie's argument from queerness. I shall treat this objection as a fourth challenge, which faces specifically expressivist accounts of the meaning of ethical claims.

This will prove important because I ultimately want to defend an expressivist account of the meaning of ethical claims; however, it is a novel inferentialist form of expressivism. So with the fourth challenge on the table, I shall then proceed to work towards my own positive expressivist account of the meaning of ethical claims and argue that it meets all four challenges which emerged in the course of our discussion of the other options.

I. TRADITIONAL EMOTIVISM

Emotivism emerged on the metaethical scene from two quite distinct directions. On the one hand, in the tradition of the American pragmatists and use theorists of meaning, Stevenson (1937) suggested that moral terms were best understood in terms of their "emotive meaning". On the other hand, in the tradition of the logical positivists and verificationist theorists of meaning, Ayer defended what he called the "Emotive Theory of Value." Ayer's statement of the position is most austere and audacious: he claimed that pure ethical claims "are not in the literal sense significant, but are simply expressions of emotion which can be neither true nor false." (1936/1946: 102-103). Here's a pregnant and oft quoted passage:

If now I generalize my previous claim and say, "Stealing money is wrong," I produce a sentence which has no factual meaning—that is, expresses no proposition which can be either true or false. It is as if I had written "Stealing money!!"—where the shape and thickness of the exclamation marks show, by a suitable convention, that a special sort of moral disapproval is the feeling which is being expressed. It is clear that there is nothing here which can be true or false... (Ibid: 107)

This is why he claims that, “in saying that a certain type of action is right or wrong, I am not making any factual statement, not even a statement about my own state of mind. I am merely expressing certain moral sentiments”(Ibid.: 107). So, initially, Ayer’s version of emotivism comes to the thesis that ethical claims do not express propositions (and are thus neither true nor false) but instead merely express moral sentiments.

But there’s something slightly misleading in this way of putting the thesis, which we can expose by considering an ambiguity in the word ‘claim’. Claims can be either linguistic acts of claiming—i.e. a type of utterance—or the linguistic medium with which these acts are carried out—i.e. indicative sentences.³⁷ In asserting that ethical claims express certain moral sentiments, Ayer must be thinking of ethical claims *qua* utterances. When we engage in the linguistic act of uttering a meaningful sentence, we often purport to display some psychological state. For example, if Ben says, “It’s sunny in Tübingen,” he purports to display his belief that it’s sunny in Tübingen; or if Amanda says, “Eew, liverwurst, gross” she purports to display her dislike of liverwurst. That is, if these utterances are sincere, then the speakers have the psychological state they purport to display in making the utterance. When it comes to the utterance of indicative sentences, it’s common to think of these linguistic acts as purporting to

³⁷ Jay Rosenberg has suggested to me that there is also something misleading about using this distinction to explain Ayer’s thesis because it obscures what is an otherwise natural reading of it—viz. that in uttering an ethical sentence one is not engaging in the linguistic act of *claiming* but in the linguistic act of *expressing*. Although I recognize that metaethicists and philosophers of language have sometimes made a distinction between these two sorts of linguistic acts—more common, perhaps, is the distinction between the act of *asserting* and the act of expressing—I think this leads to the confusion of thinking that the mere expressive character of a linguistic act prevents it from being (also?) an act of claiming or asserting. Here I propose to initially follow recent usage in the metaethical literature and treat ‘claim’ as ambiguous between linguistic act and medium and neutral with respect to the type of mental state expressed by the act or the propositionality of the medium. Likewise (although I find this more forced) I shall follow recent metaethical literature and treat ‘judgment’ as the mental correlate of a claim, which leaves open the question of whether ethical judgments are cognitive states. Some may balk at these ways of using the terms ‘claim’ and ‘judgment’; this is, I believe, a reflection of intuitive dissatisfaction with the noncognitivist *cum* quasi-realist program of recapturing ‘cognitive’ sounding language for the noncognitivist, which I shall discuss below in section II. I sympathize with the dissatisfaction and below attempt to articulate the core problem, in order to go on to develop an expressivist view that is explicitly not noncognitivist. I hope this sublimates any terminological bad faith that I engage in early on in this chapter by following recent metaethical usage of the terms ‘claim’ and ‘judgment’.

display the speaker's belief. However, Ayer's view is that the utterance of an ethical indicative sentence merely expresses moral sentiments and thus doesn't also express beliefs. In contrast, Ayer's correlative view that ethical claims do not express propositions (and thus are not true or false, i.e. have no factual meaning), turns on thinking of ethical claims *qua* sentences. Several different tokens of one orthographic sentence type or even several different tokens of appropriately modified different orthographic sentence types can, in some sense, express the same thing. For example, Matt's use of the sentence, "The strike starts tomorrow" on Wednesday and Suzy's use of the same sentence on Wednesday and Benjamin's use of the sentence, "The strike starts today" on Thursday all express the same thing. When it comes to indicative sentences, it is common to think of what is expressed by sentences as a proposition. However, Ayer's view is that, unlike descriptive sentences, which can be thought of as expressing a proposition, ethical sentences do not express a proposition.

I think we can capture this difference by distinguishing between two distinct senses of 'express'. To use narrow senses of these terms, sometimes expression is a *semantic* relation, as in "Tony's claim 'The cat is on the mat' expresses_(s) *the* thought that the cat is on the mat." Other times it is a *pragmatic* relation, as in "Tony's claim 'The cat is on the mat' expresses_(p) *his* thought that the cat is on the mat."³⁸ This corresponds to the ambiguity in the notion of a 'claim' and so is not just a difference in the character of the relation but also in the *relata*. Expression_(s) is a semantic relation that a sentence stands in to a proposition, while expression_(p) is, at least to a first approximation,³⁹ a pragmatic relation of conveying that an utterance stands in to a mental

³⁸ These roughly correlate to the senses of 'express' Sellars (1969) calls the semantic and action senses of express in accounting for the different senses in which it is appropriate to talk a claim expressing a thought. Sellars identifies a third sense of 'express' which he calls the causal sense of express. This is the sense in which smoke expresses fire and hair-color can express one's genes. Although Ayer's occasional assimilation of ethical claims to verbal ejaculations might encourage using the causal sense of express to explain expressivism, I think this is an unfruitful way to develop the theory.

³⁹ Below I argue that noncognitivists have to transmute their claims about what ethical claims express_(p) into accounts of the meaning of ethical claims, so the expressive_(p) force of ethical claims can be no *mere*

state. With this distinction in hand, we can see that Ayer's view comes to at least the following two distinct theses:

- (i) ethical claims don't express_(s) ethical propositions;
- (ii) ethical claims don't express_(p) beliefs in an ethical proposition.⁴⁰

But then what *do* ethical claims express on Ayer's view?

Because, as we saw in our discussion of internalism in chapter 2, there seems to be some important connection between ethical claims and action, noncognitivists typically argue that ethical claims express_(p) a mental state integrally caught up in motivating action. A prominent account of the psychology of human action comes from Hume. He famously argued that "...reason alone can never be a motive to any action of the will; and ... it can never oppose passion in the direction of the will"(1739: II.iii). This is why he held that "reason is a slave to the passions." According to him, there are two species of reasoning: first, demonstrative inferences about the relations of ideas and, second, probabilistically inferences about empirical matters of fact. He argues:

...it scarce will be asserted, that the first species of reasoning alone is ever the cause of any action. As its proper province is the world of ideas, and as the will always places us in that of realities, demonstration and volition seem, upon that account, to be totally remov'd, from each other.(Ibid.)

The idea is that making an analytic inference, e.g., from the fact someone is a bachelor to the fact that he is an unmarried man, cannot by itself motivate one to do something. Regarding the second species of reason, Hume allows that it is connected to volition. He writes,

pragmatic matter. Later in this chapter, I shall consider more closely the nature of this pragmatic relation and its relata in order to suggest that there's an important distinction between three different expression_(p) relations and their apparent relata.

⁴⁰ An "ethical proposition" is one that would be expressed_(s) by an ethical sentence if ethical sentences expressed_(s) propositions. It's important to add this complexity because noncognitivists can and usually do grant that ethical claims typically express_(p) a whole complex of mental states which can include beliefs in nonethical propositions. Ayer's classic example was the claim "You acted wrongly in stealing that money", which indisputably expresses_(p) the belief that you stole that money; but, according to the noncognitivist, does not express_(p) the claimer's belief in the ethical proposition that this act was wrong.

... when we have the prospect of pain or pleasure from any object, we feel a consequent emotion of aversion or propensity, and are carry'd to avoid or embrace what will give us this uneasiness or satisfaction. 'Tis also obvious, that this emotion rests not here, but making us case our view on every side, comprehends whatever objects are connected with its original one by the relation of cause and effect. Here then reasoning takes place to discover this relation and according as our reasoning varied, our actions receive a subsequent variation.(Ibid.)

But, as we see here, Hume thinks that this second species of reason is connected to volition only in that it helps us to figure out what causes pain and pleasure, and the resulting aversions and propensities are ultimately necessary for motivation. So, according to Hume, the second species of reasoning serves to direct the will, but only in conjunction with pain and pleasure. He writes,

... the impulse arises not from reason, but it is only directed by it. 'Tis from the prospect of pain or pleasure that the aversion or propensity arises towards any object: And these emotions extend themselves to the causes and effects of that object, as they are pointed out to us by reason and experience.(Ibid.)

One plausible and prominent way to think about the psychology of action that emerges from Hume's argument here is that it comes to the claim that there are two basic kinds of mental states: one that is capable of truth and falsity thus in the purview of reason and one that is intrinsically motivational and thus not in the purview of reason—or, as they are commonly called, cognitive and conative mental states.⁴¹ If we accept this distinction, we can start to see why traditional noncognitivists like Ayer also endorse the following positive thesis

(iii) ethical claims express_(p) a conative mental state.

The idea was that ethical claims were used to display one's conative attitudes towards some action or possible action, which was thought to be part of a complex use of language whereby we influence the actions and emotions of others by displaying our own conative, i.e. motivational, mental states.

⁴¹ The exact difference between these is then often spelled out in terms of the metaphor of a difference in "direction of fit with the world". See Smith (1987). Below in section III, I shall consider this metaphor in more detail because I think there is a way to understand it that actually cuts across rather than with the cognitive-conative distinction. But here the important point is one I don't want to deny, viz. that there are motivational states that are not belief-like states.

The conjunction of (i), (ii), and (iii) represents a pretty radical view—why would anyone hold it? Ayer’s official reason has in part to do with Moore’s open question argument. Recall that Moore argues that the concept ‘good’ cannot refer to a natural property because it cannot be analyzed in nonethical terms that refer to natural properties—or, correlatively, claims about what is good cannot describe a natural fact because they cannot be analyzed into nonethical claims that describe natural facts. He argues that this implies that ethical claims either describe a nonnatural fact or are meaningless. Since, however, we understand ethical claims, they cannot be meaningless, so they must describe a nonnatural fact. Ayer suggests that the final inference in this argument turns on a false dichotomy. He writes,

We begin by admitting that the fundamental ethical concepts are unanalyzable...So far we are in agreement with the absolutists. But, unlike the absolutists, we are able to give an explanation of this fact about ethical concepts. We say that the reason why they are unanalyzable is that they are mere pseudo-concepts. The presence of an ethical symbol in a proposition adds nothing to its factual content.(Ibid.: 107)

So Ayer would agree that ethical claims are meaningless, if by “meaningless” Moore meant “have factual content”; however, Ayer thinks that we can attribute a different sort of meaning to them. He writes, “we may define the meaning of the various ethical words in terms both of the different feelings they are ordinarily taken to express, and also the different responses which they are calculated to provoke”(Ibid.: 108). And this means that even if Moore is right that ethical claims cannot be analyzed in natural terms and that they are meaningful, they need not necessarily describe nonnatural facts.

But even if Moore were to grant that there is a third option, that wouldn’t force him to take it; he could stick to his guns and maintain that ethical claims nonetheless get their meaning by describing nonnatural facts. What favors Ayer’s view over this view? Ayer seems to have had something like Mackie’s challenge in mind when he writes, “In admitting that normative ethical concepts are irreducible to empirical concepts, we seem to be leaving the way open for the ‘absolutist’ view of ethics—that is the view that statements of value are not controlled by

observation, but only by a mysterious ‘intellectual intuition’”(Ibid.: 106). That is to say that since ethical facts (were they to obtain) would not be open to empirical investigation, they are somehow suspect. But Ayer agrees with Mackie that that *sui generis* ethical facts and a correlative intuitive faculty of ascertaining them would be too queer to reasonably admit into our ontology and epistemology. Nevertheless, one can defuse Mackie’s argument from queerness by rejecting the premise that ethical claims assert the existence of objective values. And one way to do this is by construing ethical claims not as expressions of belief, which aim at truth, but as expressions of some nonbelief sentiment. I think this is a central motivation for the emotivist version of expressivism.

Thus, in terms of the three challenges to an account of the meaning of ethical claims that have emerged thus far, the emotivist version of expressivism has a very clear way to answer Mackie’s challenge—i.e. to be able to defuse the argument from queerness (see page 14 above). As we saw in chapter 2, nonreductive a posteriori naturalists can also answer Mackie’s challenge, but their answer undermines their ability to answer the action-guiding challenge—i.e. to be consistent with and preferably explain the truth of rationalism and (weak) internalism. Noncognitivists have often noted that their view fairs better than naturalism on this score in that conative mental states are conceived of as intrinsically motivational states, so if ethical claims express a sort of conative state, then there is an obvious explanation of why one who makes an ethical claim is, when normal, defeasibly disposed to the appropriate action. This means that emotivism can underwrite internalism.

I think they have a tougher time with rationalism. Why, in the noncognitivist’s view, does its being true that ϕ -ing in C is good entail that one in C has a reason to ϕ ? On the noncognitivist’s view, claiming that ϕ -ing in C is good isn’t really to express a propositionally articulated belief at all, so it is sort of mysterious how it could be “true” or entail anything at all. One way noncognitivists can respond to this worry is to argue that claiming that one in C has a reason to

ϕ is itself a noncognitive claim, and talk of “truth” and “entailment” here is just a fancy way of saying that these two claims are tightly connected in that the sort of conative state expressed by the ethical claim normally (or necessarily or should) give rise to the sort of conative state expressed by the claim about reasons. This response isn’t completely satisfactory because it makes it mysterious how reasons for acting might transcend one’s conative attitudes towards various actions; however, some noncognitivists might be willing to bite this bullet while others might try to explain the very thought (comprising the bullet) in terms of the expression of conative attitudes (say) towards being humble in our moralizing. In any case, noncognitivists seem to have a leg up on naturalists with respect to the action-guiding challenge. And recall that relativists also had a leg up on naturalists with respect to this challenge but didn’t clearly meet it in that they could either capture internalism or rationalism but not both completely. So, without getting too picky about weights and measures, perhaps it’s fair to say that relativists and noncognitivists are on even par with respect to the action-guiding challenge.

However, concerning the relativist’s challenge—i.e. to explain the existence of recalcitrant difference in ethical opinion without disrespecting the strong semantic intuitions that this difference represents a genuine disagreement—noncognitivists often claim that their account gains important advantages over metaethical relativism precisely at this point. This is because they can explain the intuition that most differences in ethical opinion are genuine disagreements rather than instances of merely deploying differently relativized concepts. The basic idea is to chalk this intuition up to the fact that the claims express agreements or disagreements in sentiments or conative attitudes, and then to argue that such disagreement is more recalcitrant due to the fact that these conative states aren’t in the purview of reason or subject to correction through perception in the same way that beliefs are. Relativists will object that “disagreement in sentiment” isn’t enough to capture genuine disagreement, and I think this is a fair worry. However, at this point in the dialectic, we should note that the emotivist does have a better explanation of the intuition that ethical disagreements are genuine. He can say that even if

disagreement in attitude isn't counted as "genuine" it is still more of an opposition between apparently conflicting ethical claims than the relativist is able to capture.

So it looks as though emotivism does relatively well on all three of our challenges. However, as many have rightly worried, the traditional noncognitivist's ability to meet the relativist challenge, comes at the cost of disrespecting other ordinary semantic intuitions. Or so I shall presently argue.

Let's start with a crude objection that I think motivates, if only implicitly, much resistance to noncognitivism. Recall that we've seen *four* sorts of accounts that one might initially adopt if one doubts the naturalistic tractability of *sui generis* ethical facts. First one might adopt a version of ethical relativism that makes ethical facts into relative facts about naturalistically tractable things; second, one might attempt to identify ethical facts with particular natural facts; third, one might adopt the emotivist view that by treating ethical claims as not fact-stating; fourth, one might follow Mackie in treating ethical claims as in widespread error also eliminates the need to posit ethical facts. We've seen problems with the first two approaches. And in chapter 2, we rejected the fourth approach out of hand because it's a Moorean fact that we all know that cutting off babies' eye-lids for fun is wrong, and since knowledge is factive this means that it's true that cutting off babies' eye-lids for fun is wrong. So any account that doesn't hold that it's true that cutting off babies' eye-lids for fun is wrong is seriously counterintuitive. However, since, according to emotivism, ethical claims are neither true nor false, the third approach seems to suffer the same *prima facie* devastating problem suffered by the fourth approach.

This is a crude objection because there's an obvious way for noncognitivists to redescribe what I've called the "Moorean fact". They can easily grant that the fact that we all know that cutting off babies' eye-lids for fun is wrong entails the fact that we're all committed to this being

wrong, and then go on to insist that this latter fact is to be given an emotivist explanation rather than a descriptivist explanation. However, less crudely (but still perhaps crudely), the underlying worry seems to be that although noncognitivists can avoid countenancing widespread error in the usage of nonembedded ethical claims, they do this by countenancing widespread error in the embedding of ethical claims in the truth and falsity predicates. The worry is highlighted in connection to the emotivist's answer to the relativist's challenge. For the emotivist argue that he doesn't have to disrespect the ordinary semantic intuitions about agreement and disagreement between ethical claims; however, it appears that he can say this only by disrespecting a different sort of semantic intuition—viz. the truth-evaluability of ethical claims. For it clearly seems to be no *semantic* mistake to assert the sentence, “It is true that *E*” for any given ethical claim *E*.

Noncognitivists may insist that even this second objection is crude. For they *can* distinguish between the untrue ethical sentences we accept and the untrue ethical sentences we reject. And because of this—they will argue—the sort of semantic intuitions they are required to disrespect are not very important. Indeed, if we grant a distinction between “loose and common” and “strict” senses of ‘true’, their error-theory vis-à-vis the truth predicate can be completely nonrevisionary. That is, their view does not imply—as relativism (and perhaps error-theory) does—that certain ordinary ways of talking about ethical sentences are radically mistaken. Rather their view implies that certain controversial philosophical theories about the nature of ethical discourse rest on a bad picture or a false analogy.

However, consideration of a third and related objection reveals that whatever sense of true is at issue, there is a problem with the noncognitivist thesis that ethical claims are not truth-evaluable. This is the so-called “Frege-Geach Problem”(Geach 1960, 1965).⁴² Geach attributed to Frege the plausible claim that “A thought may have just the same content whether you assent to

⁴² See also Ross (1939: 33-34) and Searle (1969: 136-141) for independent developments of essentially the same objection.

its truth or not; a proposition may occur in discourse now asserted, now unasserted, and yet be recognizably the same proposition”(1965: 449). Then Geach argued that a similar point applies to ethical claims. That is, there is a undeniable semantic similarity between a sentence used alone to make an ethical claim and that same sentence imbedded in some unasserted context. However, it is far from clear that the noncognitivist can explain what is semantically similar between a sentence used on its own to make an ethical claim and the same sentence embedded in some unasserted context. To motivate this point, Geach offers the following argument as an example:

(P1) If tormenting the cat is bad, getting your little brother to do it is bad.

(P2) Tormenting the cat is bad.

Ergo, getting your little brother to torment the cat is bad.(443)

This argument seems straightforwardly valid, but typically its validity would be explained by the fact that its form is *modus ponens*—i.e. if *p* then *q*, *p*, thus *q*—and arguments of this form are truth preserving. But the emotivist cannot appeal to this explanation for two reasons. First, it’s not clear that she can identify and explain some semantic similarity between “tormenting the cat is bad” as it occurs in (P1) and (P2). Typically one would say that they have the same content because they express the same proposition or have the same truth-conditions; however, the emotivist will deny that (P2) expresses a proposition or has truth-conditions. Rather she will say that it expresses some conative state, but what then does it express when its the antecedent of (P1)? Since it is there unasserted, it seems wrong to say that it expresses_(p) the same conative mental state as it expresses in (P2). Even someone who wouldn’t affirm (P2) might affirm (P1), which seems to indicate that the antecedent of (P1) doesn’t express_(p) the same mental state as (P2). Second, the emotivist holds that (P2) and the conclusion aren’t truth-evaluable at all, so she cannot say that the argument is valid because it is *truth* preserving. Are there any avenues open to the emotivist for responding to this objection? They usually take one of two approaches.

The first approach is to deny one of the crucial assumptions of Geach's objection by claiming that the correct account of semantic similarity and validity is *not* one that appeals to truth-conditions and truth-preservation respectively—perhaps it is one that instead appeals to similarities in use or assertability conditions and preservation of assertability. If that's right, then we can explain the similarity of content between (P2) and the antecedent of (P1) and the validity of the argument just as we would for any other claims of the same form, in terms of use or assertability-conditions. This approach is, however, open to rejoinders. Even if it can explain the similarity of semantic content between the antecedent of (P1) and (P2) without assuming that they are truth-evaluable, this move doesn't seem like it will work for an alternative ethical argument that writes the truth predicates directly into the content of the claims. For example,

(P1') If it is true that tormenting the cat is bad, getting your little brother to do it is bad

(P2') It's true that tormenting the cat is bad.

Ergo, getting your little brother to torment the cat is bad.

This argument seems just as valid as the previous one, but it is unclear what the noncognitivist is to say about what is expressed by (P1') and (P2'), which makes it hard to see how he will explain the validity of this argument.⁴³ Because of these problems, it is much more common for noncognitivists to take the second approach to responding to the Frege-Geach problem.

⁴³ Jay Rosenberg suggested to me that it should be clear from what I've said above what the emotivist will say—viz. "That there's a 'loose and common' sense of 'true', carrying no relational semantic or metaphysical implications, according to which (P2') and the antecedent of (P1') in essence differ only *stylistically* from (P2) and the antecedent of (P1). Appeal to assertability conditions thus validates the second version of the argument in the same way as it validates the first version." However, I think this will not do. Although clearly there is a version of the argument using the loose and common sense of 'true', for which this explanation will work, given the emotivist's own distinction between loose and strict uses of 'true', one might insist that there is also a version of the argument using the strict sense of 'true' and this version of the argument seems valid. To be sure, emotivists deny that ethical claims can, in the strict sense, legitimately be called true; however, that is beside the point here since the argument still seems valid even if both premises are false—and it's not clear how assertability conditions can explain this apparent validity.

The second approach is to responding to the Frege-Geach problem is to reject “realist” conceptions of truth (e.g. correspondence theories) in favor of minimalism about truth. Minimalism is a positive view about the scope and function of the truth-predicate combined with a negative view about the semantics of the truth-predicate. The positive view is that any declarative sentence ‘S’ is true iff S; the negative view is that the truth-predicate doesn’t denote a property (e.g. correspondence) and so calling a sentence S true doesn’t carry any more metaphysical implications than simply affirming S.⁴⁴ This means that declarative ethical sentences are indeed truth-evaluable, and so the similarity in content between (P2) and the antecedent of (P1) can be explained in terms of a sameness of truth-conditions and the argument’s validity can be explained in terms of the preservation of truth. Recall, however, that emotivism included the negative thesis (i) that that ethical claims do not express_(s) truth-evaluable content. If in response to the Frege-Geach problem we now say that ethical claims *do* express something that is truth-evaluable, then it seems that we’re giving up on this part of the traditional emotivist position.

This is not, in and of itself, a bad development as long as we keep track of what we’re talking about. But before considering the prospects of an alternative version of expressivism, I want to point out that this second approach concedes something crucial to Geach’s objection—viz. that our treatment of the truth predicate, similarity of content, and validity should be the same when considering ethical and nonethical claims. If this is right—which, officially, I want to remain neutral on—then the “loose and common sense of ‘true’” response to the first two objections is

⁴⁴ There is much debate in the theory of truth about the nature and viability of, in particular, minimalism as a theory of truth. See Blackburn and Simmons (1999). Many philosophers accept that the Tarskian T-sentences are true for all and only meaningful indicative sentences and take this, as Tarski arguably did, to be a metaphysically neutral claim. The controversy begins when we ask what implications this has for the property and or concept of truth. Some prominent minimalists (e.g. Horwich (1990)) argue that because the T-sentences exhaust the meaning of the truth-predicate, truth is not a robust property and the concept of truth can play no interesting philosophical explanatory role. Bar-On and Simmons (forthcoming) criticize this distinguish between deflationism about the property of truth and deflationism about the concept of truth and criticize the latter. However, I think it is important to notice that all the noncognitivist needs in endorsing minimalism about truth is the former.

no good. For that response turns on making a distinction between the way the truth predicate applies to ethical and nonethical claims, but this minimalist response to the Frege-Geach problem undermines such a distinction.

So, the upshot of our discussion of the three objections we've considered so far seems to be that emotivists should change their position so that they can hold that ethical claims are truth-evaluable in just the same way that non-ethical claims are, and this way should be minimalist.⁴⁵ This move involves a serious commitment, because minimalism is by no means an uncontroversial approach to truth.⁴⁶ However the advantages for the expressivist of adopting minimalism are fairly clear. It allows him to abandon the negative thesis (i) in the emotivist view, which is a source of much of the problems we've seen attaching to that view.

II. QUASI-REALISM

It's thus perhaps unsurprising that we see exactly this development among contemporary expressivists such as Blackburn (1993, 1998), Timmons (1998), and Gibbard (2003). Using language that traditional emotivists such as Ayer and Stevenson would have never used, Blackburn writes, "Expressivism claims that the ethical proposition is something that we synthesize for a purpose. Its role is to act as a focus for practical thought." And, according to him, the "ethical proposition" is properly, though minimalistically, truth-evaluable:

⁴⁵ William Lycan has suggested to me that emotivists could also allow that there are two distinct truth concepts—one minimalist and one substantivist. This would allow them to avoid commitment to (i) by holding that ethical claims are truth-evaluable in the minimalist way, while maintaining that other sorts of claims are truth-evaluable in some more substantive way. But, like Lycan, I find it *prima facie* implausible that two different concepts are expressed by the word 'true'.

⁴⁶ Minimalism here is any view of truth according to which calling a sentence true doesn't commit one to the existence of a correlative truth-maker or fact to which this sentence corresponds. Most of the proponents of minimalism are deflationists, which means that they see the concept of truth as having no explanatory role in (at least) the theory of meaning. However, I doubt that the only way to be a minimalist about truth is to be a deflationist. One might also endorse a Peircian account similar to Misak (1991) or a pragmatist-prosentential account of truth similar to Brandom (1994).

We must remember Wittgenstein's dismissive attitude to invocations of truth and representation when he is dealing with the kinds of commitment that interested him. Just *because* of minimalism about truth and representation, there is no objection to tossing them in for free, at the end. But the commitments must first be understood in other terms.(1998: 80)

And in his most recent book, Gibbard is explicitly agnostic as to whether there is anything *more* to the notion of truth, but he claims to use the minimalist sense of the word 'true' throughout the book. He writes, "As for the word 'true', I accept minimalism for the term as I use it, leaving it an open question whether there's a clear and more demanding sense of 'truth'"(2003: 18). Some will wonder whether the resulting view is still a version of "noncognitivism", given that it admit that ethical claims are truth-evaluable. As I've written elsewhere (2005: 409), I think the resulting view deserves to be called "noncognitivism" just as much as the current 1st Cavalry Division of the U.S. Army deserves to be called a "cavalry" division. And the distinction between the semantic and psychological senses of the expression relation that was operative in my initial characterization of emotivism above can help to make clear what's at stake here.

Blackburn and Gibbard abandon (i), but they retain their commitment to (ii) and (iii). That is, they concede that ethical and nonethical claims alike express_(s) propositions, but they maintain that the psychological state expressed_(p) by ethical claims is interestingly different from the psychological state expressed_(p) by descriptive claims. Specifically, given the Humean distinction between cognitive and conative states of mind, this new view can claim that while ethical claims express_(s) propositions, they also express_(p) conative states of mind rather than cognitive states of mind.⁴⁷ This contemporary version of noncognitivism which is the combination of (ii) and (iii) with minimalism about truth is what Blackburn (1984, 1993, 1998) (and following him, Gibbard(2003)) calls "quasi-realism".

⁴⁷ Initially it may seem incoherent to say that a sentence expresses_(s) a proposition, but the uttering of this sentence does express_(p) a conative mental state rather than a belief in the proposition expressed_(s) by the sentence. However, reflection on the common usage of sentences such as "I'd like tea" or "I'm going to the movies" reveals that this is not at all incoherent. These sentences can plausibly be thought to express_(s) a proposition, although when we utter them we are not typically expressing_(p) our belief in proposition expressed_(s) by the sentence but rather, respectively, a desire and an intention.

Is the quasi-realist account of the meaning of ethical claims, which results from dropping (i) while retaining (ii) and (iii), a plausible view? Arguably it has the same resources to meet the relativist's challenge as its emotivist cousin. However, I think it faces problems deriving from its ability to meet the action-guiding challenge and Mackie's challenge. Let's take these two lines of criticism in turn.

We can bring the first problem into view by considering how exactly noncognitivists (of any stripe) avoids the main problem with metaethical relativism, which was that relativism has to disrespect ordinary semantic intuitions about the dialectical connections between some ethical claims. The noncognitivist argues that a disagreement in conative mental state is enough to explain our intuition that ethical claims such as,

(1) Taking the money is wrong,

and

(2) Taking the money is not wrong,

stand in genuine disagreement. However, even if we grant that there are other sorts of disagreement than factual disagreement, it's questionable whether a mere disagreement in conative state can capture the sort of disagreement we sense in (1) and (2). To see this, consider a case which is more clearly a mere disagreement in conative state:

(1') Coconut ice-cream is delicious.

(2') Coconut ice-cream is disgusting.

Here there seems to be a clear sense in which each "disputant" could recognize that although she has a different preference than the other "disputant" that doesn't imply anything about how the other should behave. With (1) and (2), however, it seems that each disputant is making a claim

that has implications about how the other person should think and behave. That is to say that from (1) it seems that we can infer

(3) We should not allow taking of the money,

and from (2) it seems that we can infer

(4) We should allow taking of the money,

and (3) and (4) stand in a sort of disagreement that isn't explained merely by the claim that they express different conative mental states. For if it were, then we could also infer

(3') We should eat coconut ice-cream,

from (1') and

(4') We should not eat coconut ice-cream,

from (2'), but clearly these inferences don't follow. This seems to indicate that ethical claims, unlike claims that merely express conative mental states, carry a more universal implication that isn't clearly captured by noncognitivism of either the traditional or quasi-realistic varieties. Hare captures this worry well:

...to say that imperatives express wishes may lead the unwary to suppose that what happens when we see one, is this: we have welling up inside us a kind of longing, to which, when the pressure gets too great for us to bear, we give vent by saying an imperative sentence. Such an interpretation, when applied to such sentences as 'Supply and fit to door mortise dead latch and plastic knob furniture', is unpalatable. And it would seem that value-judgments also may ...indeed be in some sense, like imperatives, prescriptive, without having this sort of thing said about them.(1952: 10)

I think the worry here is indicative of a deeper worry about the fact that conative mental states don't seem to be the objects of rational deliberation in a way necessary to make sense of the role of rational deliberation in ethical discourse or, ultimately, the truth of rationalism. This isn't to say that we cannot have reasons for having particular desires or feelings, but the reasons don't attach to the content of conative states in the same way as they do to the content of cognitive states. If ethical claims express some mental state that isn't subject to reasoning in the

same way as beliefs or cognitive states, then we may worry that ethical thought and discourse operate more like coercion, manipulation, and admonishment than like reason giving. However, despite the undeniable fact that ethical claims also share something with these things that don't directly provide reasons for actions, it seems that any account that cannot explain the way that ethical claims can be taken up into practical reasoning both as premises and conclusions of genuine ethical reasoning is missing something crucial to ethical thought and discourse.

Arguably what is missing is a plausible explanation of the truth of rationalism. Recall that rationalism is (in one version) the view that if it's true that ϕ -ing in circumstances C is good, then anyone in C has a reason to ϕ . Quasi-realists have no problem *endorsing* this view, for they have no problem with the truth-evaluability of the ethical judgment or capturing a sense of implication on which this implication appears cogent. However, it's hard to see that they have a plausible *explanation* of rationalism in the offing. Rationalism seems designed to capture the way that ethical discourse is a form of practical (justificatory) *reason giving*. Ethical truths have implications for how we ought to act, because of this ethical claims seem potentially open for justification and rational defense. That is, they both provide (justificatory) reasons for action and call out for (justificatory) reasons for endorsement. By treating ethical claims as expressive of conative mental states, the noncognitivist seems to obscure precisely this rationalist feature of ethical discourse. This is because neither conative states nor their expressions stand in genuine justificatory relations.

To be clear, I don't take this to be a knock-down argument against quasi-realism. For there are ways for quasi-realism to save appearances. For instance, they can maintain that the doctrine of rationalism is merely the expression of a de facto psychologically causal connection between two types of mental states. One who has the conative state expressed by the ethical claim will be caused (perhaps contingent on being normal and having the necessary concepts) to have the mental state expressed by the claim about reasons. I doubt that this is true, but there is

enough wiggle room in the notion of “type of mental state” to make it hard to refute. Perhaps a more plausible quasi-realist rejoinder is to argue that the statement of rationalism itself is a claim that requires a second-order quasi-realist analysis, for claims about the justificatory relationship between claims about what is good and claims about reasons are themselves normative claims requiring quasi-realist analysis. This also isn’t a completely satisfactory response for at least two reasons. First, it puts serious pressure on the concept of a “conative” mental state: what exactly is the nature of the conative state expressed by such rarified normative claims, and is there independent psychological reason to believe that there are such nonbelief states? Second, the rejoinder threatens to make too much noncognitive: are all claims about justification noncognitive, and if so does that mean that all claims about knowledge or even all claims (period) are noncognitive? Again, there is wiggle room here, but there is also reason to be skeptical of noncognitivism.

Now let’s turn to the question of whether quasi-realism can answer Mackie’s challenge while overcoming the problems besetting emotivism. Notice that the crude objections put to emotivist versions of expressivism above can be easily recast in psychological terms in order to be put to quasi-realist versions of expressivism. For example, emotivism avoids an error theory about unembedded ethical claims only by treating the ordinary use of the truth and falsity predicates in conjunction with ethical claims as, strictly speaking, erroneous. And a very similar worry threatens quasi-realism, for it too must treat the embedding of ethical sentences in belief contexts, e.g., “I believe that stealing is wrong”, as, strictly speaking, erroneous. Relatedly, we worried that emotivism has to countenance a strange sort of semantic ignorance about the proper application of the truth predicate; and again, it seems that the parallel point applies to quasi-realism with respect to the application of belief-contexts.

Typically the quasi-realist response to such crude objections is to grant the point that there’s a sense in which we call the mental state expressed by ethical claims and descriptive claims alike

a “belief” but to insist that there is also a more strict technical sense in which we can distinguish descriptive beliefs from ethical beliefs. Before, appeal to multiple senses of the offending terms appeared to not save emotivism because the most attractive response to the Frege-Geach problem required treating ‘true’ univocally. Interestingly, the Frege-Geach problem has a much harder time undermining quasi-realism⁴⁸ (after all the position is designed to avoid the Frege-Geach problem), which means that there is more room for the quasi-realist to respond to these crude objections by arguing that there are “loose and common” and “technical” senses of ‘belief’. And, while both are perfectly legitimate, when he claims that ethical claims express_(p) conative states rather than beliefs, he has the technical sense in mind.

Whether this is a good response to the objections depends on how well the quasi-realist can specify the distinction between the two senses of ‘belief’, but on the face of it this seems like it will be a tricky project since the loose sense of ‘belief’ will have to cover both cognitive and conative—i.e. *noncognitive*—mental states. Moreover, this problem seems to spread to other propositional attitudes that are conceptually related to belief. It seems, for example, just as semantically in order to say, “I believe that genocide is wrong” as to say, “We all know that genocide is wrong.” Typically, however, claims of the form ‘S knows-that p’ are thought to be true only if S has a belief that p. However, if the claim, “Genocide is wrong” expresses_(p) a noncognitive mental state, then it is puzzling how the quasi-realist can count this as the mental component of this knowledge.

I suspect that these problems are merely symptomatic of a deeper problem. To bring this problem into view, recall that expressivism is supposed to be an account of the meaning of ethical claims that has the resources to provide a nondescriptivist answer to Mackie’s challenge.

⁴⁸ There is actually a whole literature about this. Some of those who think that quasi-realism still has a problem with the Frege-Geach problem Hale (1993), Dreier (1996), Zangwill (1990); however, see Horwich (1993), Blackburn (1993), Blackburn (1998), and Gibbard (2003) for some responses. In the following section, I’ll return to the Frege-Geach problem and argue that, while it doesn’t clearly undermine quasi-realism, it points the way to a deeper problem.

Emotivists win their answer to Mackie right upfront by denying that ethical sentences are truth-evaluable, i.e. expressive of propositions; then they fill in a semantic account in the wake of this denial. Quasi-realists want to avoid the problems generated by the first move by endorsing minimalism, but then what's their answer to Mackie's challenge? It must be that although there's a (minimalist) sense in which ethical claims are truth-evaluable, the interesting differences between ethical and descriptive claims comes at the psycho-pragmatic level of what sort of mental state is expressed by each sort of claim—here they follow in the spirit of even if not letter of emotivism. However, if this is to continue to offer an answer to Mackie's challenge, then the difference between what is expressed by ethical and descriptive claims cannot be a solely pragmatic matter. Many sorts of statements can be differentiated pragmatically in terms of what type of psychological state they express, without bearing on the issue of whether they embody a commitment to some sort of fact. For example, first-personal avowals of emotions—e.g. “I love you”—express the relevant emotion, and simple requests—e.g. “I'd like eggs”—express desires, but that does not imply that we should treat them semantically as nonfactual.⁴⁹ So for contemporary quasi-realist versions of expressivism to continue to offer a genuine response to Mackie's challenge, they need to treat the *psychological* difference in what kind of mental state ethical and descriptive claims can be used to express as making the crucial semantic difference between the meaning of ethical and descriptive claims.⁵⁰

⁴⁹ Of course, on some deflationary accounts of self-knowledge such avowals *are* treated as non-fact-stating, but I think these accounts are making the same mistake as would be made by a noncognitivist who inferred from the fact that ethical statements express_(p) attitudes that ethical claims are nondescriptive. Not all expressions_(p) of attitudes are appropriately thought of as nondescriptive. See Bar-On (2004: 226-284) for more discussion of simple expressivist accounts of avowals and argument against this inference.

⁵⁰ Thus, I think it is no coincidence that Gibbard (2003) starts with a distinction in terms of pragmatic force between factual and plan-laden statements and then transmutes this into a distinction about the type of content had by each sort of statement. He writes, “In a different and more traditional sense, one could treat belief that I am about to pack and a decision to pack as having the same ‘content’, *that I will forthwith pack*, toward which I take different propositional attitudes: belief, and deciding to actualize. The two attitudes have different “directions of fit”. we can say, towards the same item of ‘content’. Here, though, I am opting to transmute force into content: I speak of a single attitude ‘accepting’ that one can take towards distinct items of content, *that I will forthwith pack* and the plan *to pack*”(2003: 47) He goes

Dreier (2004) has argued that the quasi-realist's ability to make out this difference is threatened by the endorsement of minimalism. For, as Dreier sees things, "Minimalism sucks the substance out of heavy-duty metaphysical concepts. If successful, it can help [quasi-realism] recapture the ordinary realist language of ethics. But in so doing it also threatens to make [quasi-realism] indistinguishable from realism." (Ibid.: 26) Thus, he asks what the real difference is between quasi-realism and realism. Obviously the difference cannot be the old difference about the truth-evaluability between emotivists and realists—for quasi-realists are distinguished from emotivists precisely in that they accept that ethical claims are truth-evaluable. Initially, quasi-realists will claim that the difference comes in their view that ethical claims express_(p) nonbelief states? Indeed, this might have been the core of the emotivist theory all along; for, as many have pointed out, given Ayer's own proto-minimalist account of truth⁵¹, he should have held that ethical sentences are truth-apt but that a commitment to their legitimacy is consistent with a denial of ethical facts because pure ethical statements are used to express mere *pseudo-beliefs*. However, Dreier argues that the problem with this suggestion is that many of the same reasons that forced quasi-realists to adopt the minimalist approach towards truth apply *mutatis mutandis* to belief. For, as we've already seen, it seems to be part of the surface phenomena of ordinary ethical discourse that speakers will freely embed ethical claims in 'belief'-contexts. Moreover, it seems natural to talk of ethical "assertions", but the notion of an assertion is hard to divorce from the notion of expressing a belief.

Timmons (1998) has recently suggested that although a quasi-realist should grant that ethical sentences are truth-evaluable, assertive when uttered, and thus expressive of beliefs, he

on to claim that his idiosyncratic use of 'content' is optional, but what is not optional for the expressivist is that the pragmatic-psychological distinction between ethical and descriptive statements make the right sort of difference in one's account of the meaning of those statements.

⁵¹ Compare Ayer: "Reverting to the analysis of truth, we find that in all sentences of the form 'p is true,' the phrase 'is true' is logically superfluous...we conclude, then, that there is no problem of truth as it is ordinarily conceived. The traditional conception of truth as a 'real quality' or 'real relation' is due, like most philosophical mistakes, to a failure to analyse sentences correctly" (1936/1946: 88-89; quoted in Dreier 2004: 24).

can distinguish himself from a descriptivist by arguing that ethical sentences are *nondescriptive* and thus not expressive of *representational* beliefs. I am skeptical of the coherence of this suggestion; however, the important point here is that Dreier's argument again applies: many of the same reasons that forced noncognitivists to adopt the minimalist approach towards truth apply *mutatis mutandis* to the notions of description and representation. Ordinary ethical discourse embeds base-level ethical sentences in the putatively realist language of 'represent' and 'describe' in a way that belies the Timmons' view that ethical claims are nondescriptive or not expressive of representational mental states. For example, one might talk about the way a character in a play was represented as evil through various descriptions of his moral nature. Perhaps this way of talking is less pervasive and more easily dispensable, but if it is really important to the plausibility of quasi-realist noncognitivism that it capture the apparently realist-sounding surface phenomena of moral discourse, then it appears that quasi-realists should adopt minimalist approaches to these alternative ways of making out the necessary distinction. That is to say that pressure to accommodate the surface phenomena of ordinary moral discourse, which has encouraged noncognitivists to adopt minimalism about truth and belief continues to put pressure on quasi-realistic versions of noncognitivism that seek to mark out the relevant distinction with things like 'representational mental state' or 'description'. This is what Dreier calls "the problem of creeping minimalism":

I insist that moral sentences are descriptions, not as a matter of metaphysics but just as a matter of ordinary language. If someone asked you to describe the character of Macbeth, you might start by saying that Macbeth is a good man turned bad by ambition...[likewise] there is plainly an ordinary sense of 'representation' that applies full well to moral beliefs. Once minimalism gets creeping, it's hard to see how to stop it.(Ibid: 29)

But is it really? As I understand it, the problem of creeping minimalism turns on the fact that each of these noncognitivist suggestions appeal to terms that are deployed as part of the surface phenomena of ordinary ethical discourse. And as long as noncognitivists continue to appeal to such ordinary terms, creeping minimalism will threaten to undercut whatever distinctions they

come up with. But what about making the distinction in terms that clearly are *not* deployed as part of the surface phenomena of ordinary moral discourse? It seems that there would then be no pressure to let minimalism undermine the distinction. In fact, this seems to be precisely the strategy pursued by Gibbard (2003). He stipulates technical meanings for the phrases “the thing to do” and “plan-laden”, writing:

Suppose, let me stipulate, the phrase works like this: to conclude, say, that fleeing the building is *the thing to do* just is to conclude what to do, to settle on fleeing the building. By sheer stipulation then, the meaning of this phrase ‘the thing to do’ is explained expressivistically: if I assert “Fleeing is the thing to do”, I thereby express a state of mind, deciding to flee.(Ibid.: 8)

And with the help of minimalism, this state of mind ultimately comes to be called a “belief” in Gibbard’s theory—for one way to say that one thinks that fleeing is the thing to do is to say that he believes that fleeing is the thing to do-- but it is a “plan-laden” belief which distinguishes it from a “prosaically factual” belief.(Ibid.: 7, 185-188) Thus, in Gibbard’s view, the crucial semantic difference between ethical and descriptive claims is the *type* of belief they express—plan-laden or prosaically factual. Initially, it is hard to see how minimalism could creep in and undermine this way of making out the distinction, since it is unclear what it would even mean to take up a minimalist approach towards the distinction between plan-laden and prosaically factual beliefs. I guess one could insist that all beliefs are prosaically factual, but there would be no pressure from ordinary ethical discourse to do so.

Granting that, however, one might reasonably wonder whether Gibbard’s view—with its adoption of minimalism and psychological distinction made in technical terms—continues to offer a recognizably expressivist answer to Mackie’s argument from queerness. How, that is, does the view that ethical claims express plan-laden beliefs make out room for (a) affirming the legitimacy of ethical discourse, (b) denying the existence of nonnatural facts with a *sui generis*

prescriptive nature, and (c) denying that ethical claims state either relative or nonrelative natural facts? To put the objection boldly, it is unclear whether there is any connection at all between Gibbard's psychological account of ethical beliefs and the epistemological and ontological status of ethical discourse. To see this, consider the debate over the epistemological and ontological status of mathematics. Here, one might reasonably think it highly implausible to treat mathematical beliefs as "plan-laden", but it would nonetheless be consistent to affirm the legitimacy of mathematical discourse while rejecting the existence of both nonnatural mathematical facts and the reduction of putative mathematical facts to natural facts—both formalism and constructivism represent such a combination. So it's nothing about plan-ladenness *per se* that implies that certain beliefs lack ontological commitment. Thus, the problem for Gibbard's view is that it is not enough just to characterize the psychological nature of ethical beliefs and then just say that the beliefs expressed by descriptive statements are different; one needs to say why this difference makes a difference to the epistemological and ontological status of ethical discourse.

Dreier proposes a tentative answer to this problem on Gibbard's behalf, which he calls the "'explanation' explanation". To this end, he draws on Gibbard's suggestion⁵² that what is distinctive about quasi-realism is that,

we can explain belief in [normative facts]...without helping ourselves to normative facts at the outset...This would contrast with a standard realist's mode of explanation—appropriate in my view, for naturalistic thoughts. To explain belief in natural fact adequately, we must assume a natural world of which we are a part. We must start with a realm of natural facts.(2003: 187)

Dreier takes the crucial point here to lie in what Gibbard means by "explanation". Dreier considers how we might explain someone's ordinary belief that the moon is a quarter of a million miles away from the earth and suggests that all sorts of things might constitute the causal etiology of the belief, which means that "the best explanation for your belief that the

⁵² as well as on suggestions from O'Leary-Hawthorne and Price (1996), and Fine (2001).

moon is a quarter of a million miles away need not involve the moon itself, nor the fact that it is a quarter of a million miles away”(2004: 38). So Gibbard cannot mean that the causal etiology of one’s naturalistic belief that *p* necessarily involves the naturalistic fact that *p* in a way that one’s ethical belief that *e* does not—for neither necessarily involve the fact believed. Dreier considers the possibility that by “explanation” Gibbard means something like a “just-so” story about how we have come to have the capacity to have the relevant sort of belief. To my mind, this is a plausible interpretation of what Gibbard means; however, it puts the quasi-realism versus real-realism debate on shaky ground—for it is not completely clear what the criteria of adequacy are for a “just-so” story or if there is a fact of the matter about which of several competing “just-so” stories are correct. Although Dreier doesn’t raise this worry (or any other) about this interpretation of what Gibbard means, perhaps it is what leads him to write, “But I think that’s not what Gibbard means...The explanation of your belief that the moon is a quarter of a million miles away, in this context, is that in virtue of which it is true to say of you that you believe it. And Gibbard is speculating that an account of what naturalistic belief consists in will appeal to naturalistic facts”(Ibid.), and to make the contrast explicit (I think) he could have added: “but an account of what ethical belief consists in will not appeal to ethical facts.” This way of explaining the distinction between quasi-realism and real-realism is what Dreier calls the “‘explanation’ explanation”.

Initially it strikes me as crazy. Why saddle the realist with the view that, for instance, an adequate account of that in virtue of which it is true to say of someone that he believes that murder is good will appeal to the fact that murder is good? But maybe there is a way around this problem. To see this, notice that there are at least two initial ways to interpret Dreier’s “‘explanation’ explanation” depending on whether the operative claim is about belief tokens or belief types. The claim could be: (a) a realist about a statement expressing the belief that *p* holds that an adequate account of that in virtue of which it is true to say of someone that she believes that *p* will appeal to the fact that *p*, while an quasi-realist denies the necessity of appealing to the

fact that p. Or the claim could be: (b) a realist about T-discourse (e.g. ethical discourse) holds that an adequate account of that in virtue of which it is true to say of someone that she has T-beliefs will appeal to T-facts, while a quasi-realist denies the necessity of appealing to T-facts.

At one point Dreier seems to mean (a). He considers the question of whether Gibbard's realism about naturalistic beliefs comes to the view that an account of what the beliefs consist in will appeal "to the very facts that are also the objects of beliefs?" He admits, "Gibbard is not specific"(2004: 44). But, appealing to Fine's (2001) proposal that realists about an area of discourse are to be distinguished from quasi-realists not by what they say about whether a simple statement S in the area relates the speaker to the fact that-S but by what they say about whether "protected" statements, such as "Tom believes that-S" relate Tom to the fact that-S, Dreier writes: "...but I don't see why he would resist Fine's suggestion here"(Ibid).

However, Fine's suggestion seems to be a variant of (a), and there's a very good reason for Gibbard to resist it: It would force us to be quasi-realists about all statements that express false beliefs. For no adequate explanation of what constitutes someone's false belief that p will appeal to the fact that p, because it's not really a fact. This was effectively the reading I assumed in expressing my bewilderment at the beginning of the previous paragraph. However, notice that (b) has a similar problem, which we can see by substituting 'false' in for T. Surely we shouldn't end up being quasi-realists about a set of false statements just because we think that an adequate account of what constitutes the false beliefs that {p₁...p_n} which they express shouldn't appeal to the facts that {p₁...p_n}—for both realists and quasi-realists should agree that *merely* putative facts should not be part of an adequate explanation of beliefs in them.

What we seem to need in order to make the "explanation" explanation work is a way to distinguish realist "explanations" from quasi-realist "explanations", which doesn't force the realist to assume that the belief to be explained is true. To this end, consider the following alternative interpretation of the "explanation" explanation: (c) a realist about a statement

expressing a belief that *p* holds that that an adequate account of that in virtue of which it is true to say of someone that they believe that *p* appeals to the fact that the believer represents *p* as being a fact, while a quasi-realist denies the necessity of appealing to this fact about the believer's representations. This avoids the problem for (a) and (b) posed by false beliefs.⁵³ And if (c) represents a reasonably generous way to interpret Dreier's proposal, Dreier's proposed interpretation of the block quote above from Gibbard comes to the following view: In order to explain what constitutes a prosaically factual belief that-*p* one must appeal to the fact that the believer represents *p* as being a fact, but this is not the case when it comes to explaining what constitutes a plan-laden belief that *p*, where one must only appeal to the believer's decision to engage in a certain course of action.

But I wonder how different this is from Timmons's attempt to distinguish quasi-realist expressivism from realism in terms of whether one thinks that the statements in question express representational mental states. For all Timmons seemed to mean in claiming that ethical statements express nonrepresentational mental-states is that the mental-state expressed by the ethical statement '*S*' is not constituted by the fact that the person represents *S* as being a fact. So, even on Dreier's proposal, it seems that in order to link the psychological account of ethical beliefs to an answer to Mackie's argument from queerness, we would have to appeal to a distinction made in terms that are, again, part of the surface phenomena of ordinary ethical

⁵³ Moreover, it's (almost) in line with the following quote from Dreier: "a realist about the moon should say that belief that the moon is a quarter of a million miles away consists in some fact that includes the moon, presumably the fact of standing in some representation relation to the moon"(2004: 44). This suggestion has two problems. First, surely not any representation of the moon will do; if I represent the moon as being made of cheese, this needn't be involved in an account of what my belief that the moon is a quarter of a million miles away consists in. Rather the obvious suggestion is that the moon-realist will hold that my belief that the moon is a quarter of a million miles away consists in my standing in the representation relation to the fact that the moon is a quarter of a million miles away. But, second, what if I believe that the moon is ten miles away—is the realist then forced to say that what this belief consists in is my standing in the representation relation to the fact that the belief is ten miles away? Hopefully not, since that is not a fact. Rather, the obvious suggestion is that the moon-realist will hold that my belief that the moon is ten miles away consists in my representing the moon as being ten miles away as a fact. I take this to be an implication of (c).

discourse, which means that the position will still fall prey to the problem of creeping minimalism.

The upshot of these reflections on quasi-realism seems to be that abandoning the emotivist's commitment to (i) is a necessary adaptation for the plausibility of nondescriptivist views. However, the quasi-realist view which results from adopting a minimalist stance towards truth and simply dropping (i) doesn't clearly have anything more to offer than emotivism with respect to the action-guiding challenge. More specifically, this means that quasi-realism doesn't clearly help capture and explain the truth of rationalism. Moreover, we learned that, once he has adopted minimalism, ordinary psychological distinctions such as belief/nonbelief or representation/nonrepresentation will not help the quasi-realist to distinguish himself from the realist because of creeping minimalism. Moreover, Gibbard's technical psychological distinction between prosaically factual and plan-laden belief will not help because it is either conceptually independent of the question of realism or it is related in a way that makes it ultimately vulnerable to creeping minimalism. Now, we might try some other technical psychological distinction, but I doubt that other technical psychological distinctions will fare better than Gibbard's. Quasi-realists need to identify an expressive difference between ethical and descriptive statements that lines up with the ontological and epistemological difference they seek to use in their answer to Mackie's argument from queerness and is not subject to the problem of creeping minimalism. Since I want to pursue an alternative nondescriptivist account in the rest of this chapter, I suggest that we use this problem to generate a new challenge to developing account of the meaning of ethical claims. The *challenge of creeping minimalism and valid reasoning* is directed only at accounts which deny that ethical claims aim to describe some independent reality. The challenge is to avoid disrespecting the intuition that ethical claims are truth-evaluable but nonetheless explain how a the account of the meaning of ethical claims

bears on ethical ontology in such a way to provide a nondescriptivist response to Mackie's challenge and the action-guiding challenge.

III. INFERENCEALIST EXPRESSIVISM

So far, we've explored the three major families of descriptivist accounts of the meaning of ethical claims and the two major families of nondescriptivist accounts of the meaning of ethical claims. Four challenges have emerged, and we've yet to find an account that satisfactorily meets all of them. The fact that none of the prominent descriptivist or nondescriptivist accounts meets all four challenges may indicate that they are too demanding. However, I think there is a pregnant corner of logical space that has yet to be adequately explored, and so I hold out hope for an account that meets all four challenges.

To begin bring the rough contours of this alternative corner of logical space into view, I want to consider again the idea that the meaning of ethical claims can be captured in terms of what they express. The debate between descriptivists and emotivists can be cast as one over the question of whether ethical claims considered as sentences $\text{express}_{(s)}$ a proposition. However, in the discussion of quasi-realist noncognitivism, we saw that there is strong reason both to think that ethical sentences do $\text{express}_{(s)}$ propositions and to think that this doesn't settle the question of whether ethical claims exhibit an ontologically and epistemologically relevant expressive contrast to descriptive claims. This is because one might want to draw an expressivist contrast not at the level of $\text{expression}_{(s)}$ but at the level of $\text{expression}_{(p)}$. $\text{Expression}_{(p)}$ is initially a pragmatic relation between ethical claims, considered as acts of utterance, and whatever these utterances show or make public. A natural way to understand the quasi-realist position is as insisting on an expressive contrast between ethical and descriptive claims only at the level of

what they express_(p). Till now, however, I was fairly vague about what it means for a claim to express_(p)—although we saw that for an expressivist it cannot be a *merely* pragmatic matter because it is supposed to serve in the account of the linguistic meaning of ethical claims. Now it is time to consider the question of how exactly an expressivist should understand expression_(p).

In the previous chapter, we saw that emotivism comes to a triad of related but different theses:

- (i) Ethical sentences don't express_(s) a proposition
- (ii) Ethical utterances don't express_(p) beliefs
- (iii) Ethical utterances do express_(p) a conative mental state

It is (i) that gets traditional noncognitivists into the most trouble, which is why contemporary noncognitivists such as Blackburn and Gibbard have dropped it in the course of adopting a minimalist approach to truth and propositionality and becoming “quasi-realists”. This means that the force of their expressivism rests on their use of expression_(p) rather than expression_(s). As we saw, however, theirs can be no merely pragmatic thesis—they must be claiming that the issue of what ethical claims express_(p) determines what ethical claims mean. There seem to be roughly two different ways this idea has been understood in contemporary metaethical debate, and now I'd like to show that both of them are ill-suited for generating defensible versions of expressivism.

The first way expression_(p) is sometimes understood is as a purely causal relation. For example, certain sorts of spontaneous cries or grimaces are sometimes thought to express pain in that they are caused by the pain. And it is because these sorts of cries and grimaces are reliably caused by pain that producing them conveys the pain. These cries and grimaces, that is, are causally indicative of pain, and, insofar as an audience is aware of this, such a cry or grimace will convey the pain. This is presumably an evolved and natural signaling system by which

important aspects of one's mental life are conveyed to conspecifics who can use that information in attempts at cooperation and coordination. Given the Humean distinction between cognitive and conative mental states that we have been operating with, the most natural way to use the model of cries and grimaces to interpret (ii) and (iii) is to claim that ethical claims—like these cries and grimaces—are also part of a (perhaps much more complex) natural signaling mechanism that depends on a reliable causal connection between the presence of a conative mental state and the making of an ethical claim.

Although some of the rhetoric of early noncognitivists encourages thinking of the position along these lines, I think this interpretation of (ii) and (iii) is radically uncharitable. For cries and grimaces are relatively simple and blunt expressive vehicles; and, as such, it is far from clear that they even have linguistic meaning. By contrast, ethical claims are patently more complex in an obviously semantical way. And, as I stressed in the previous chapter, expressivism makes contact with Mackie's challenge only insofar as it offers an alternative account of the linguistic meaning of ethical claims. And, even if ethical claims are causally indicative of some conative mental state, giving this as an account of their meaning doesn't do justice to their semantical complexity.

Moreover, the expressivist wants to say that ethical claims $\text{express}_{(p)}$ something different from descriptive claims, even when they are made insincerely (notice that (ii) and (iii) aren't conditional on sincerity). But thinking of $\text{expression}_{(p)}$ as a purely causal relation cannot make sense of the possibility of insincere ethical claims. For part of what it is for a claim to be insincere is for it to be made in the absence of the mental state it purports to convey, but, if this is true, then this mental state cannot be what causes the making of the claim. To be sure, one can fake a cry or grimace, and so one can use these expressive vehicles to purport to convey pain even when one is not in pain. However, understanding how this is done requires something more than a causal explanation. One needs an account of the way in which it is possible to

intentionally follow the conventions attaching to particular expressive vehicles to dissimulate particular mental states.

Because of these problems, I think it is best to set aside the purely causal understanding of $\text{expression}_{(p)}$ in debate about expressivism. A more charitable way to interpret (ii) and (iii) is by understanding $\text{expression}_{(p)}$ as a conventional relation (which might, of course, supervene on some highly complex causal relations). It is a common idea that utterances show or display how their author is thinking or feeling about something, but clearly they usually do this not by being mere causal indicators but by conforming to violable expressive conventions. For example, by exclaiming “Hooray” or clapping one’s hands, one can show enthusiasm because of certain linguistic and gestural conventions surrounding these expressive acts. To use this idea to interpret (ii) and (iii), one would have to argue that the linguistic conventions surrounding ethical claims are such that these claims convey not a belief but a conative mental state.

This conventional understanding of $\text{expression}_{(p)}$ easily overcomes both of the problems with the causal understanding. Expressive conventions add a level of complexity that can serve in a plausible account of the patent semantical complexity of ethical discourse. Moreover, with expressive conventions in play, it is easy to understand the way in which insincere ethical claims might be thought to express conative mental states. As long as one can, in general, put on the pretense of conveying some mental state by conforming to the relevant expressive conventions even when one doesn’t possess the mental state, then there should be no special difficulty with doing this for conative mental states. Naturally, there will be pressure towards sincerity from the fact that conventions break down when they are consistently flouted. Clearly, however, this doesn’t undermine the possibility of insincerity.

Defending any noncognitivist version of expressivism on this conventional understanding of $\text{expression}_{(p)}$ requires defending the view that there are indeed expressive conventions responsible for the right sort of expressive contrast between ethical and descriptive claims. What

is the “right sort” of expressive contrast? Recall that, in order to meet Mackie’s challenge, the expressivist story needs for there to be a contrast between what is expressed by ethical and descriptive claims that explains why ethical claims, by contrast to descriptive claims, do not carry an ontological commitment to a realm of facts and an epistemological commitment to the in principle knowability of these facts. It’s clear enough how noncognitivists want to use expression_(p) to do this. They want to say that the expressive conventions are such that, while descriptive claims are conventional means of conveying a cognitive mental state (a belief), ethical claims are conventional means of conveying a conative mental state (a moral attitude); and the difference between cognitive and conative mental states makes the relevant epistemological and ontological difference. But what is the argument for the view that the expressive conventions are this way?

It sometimes seems as if noncognitivists think that it is enough to point out that many ethical claims convey a motivation to act in a particular way, and, given a Humean account of motivation, this means that ethical claims convey noncognitive mental states. However, even if we accept the two premises of this argument, it is not powerful enough to establish the sort of expressive contrast between ethical and descriptive claims needed for expressivism. For, plenty of descriptive claims also convey a motivation to act. For example, one who says, “I’m really hungry” has presumably conveyed some motivation to seek out food; and, given the Humean account of motivation, this would mean that he has conveyed a noncognitive state. But we wouldn’t want to say that this means that his claim about his is not descriptive.

This indicates that it is incumbent on the noncognitivist to argue not only that the expressive conventions surrounding ethical claims are such that these claims convey a noncognitive mental state but also that these claims do *not* convey a cognitive mental state. On the face of things, however, this will be a tough case to make. For it seems that even the expressivist should admit that the mental states expressed by ethical claims exhibit the logical-semantic properties

typically thought to be characteristic of cognitive mental states. After all, the expressivist's project is to explain the meaning of ethical claims in terms of what they express; so whatever the account is of what they express is, it will need to be able to explain the obvious semantic-logical features of ethical discourse. And, as we saw above, Geach (1965) famously argued that, by trying to explain the meaning of ethical claims not in terms of their truth-conditions but in terms of the noncognitive mental state allegedly expressed, noncognitivists seem to lose the ability to explain the intuitive sense in which ethical claims exhibit a strong semantic similarity in asserted and unasserted contexts. So any plausible version of expressivism needs an answer to this objection. Yet the noncognitivist view that we've articulated so far—viz., that the expressive conventions surrounding ethical discourse are such that ethical claims convey not cognitive mental states but noncognitive mental states—seems to be directly at odds with answering the Frege-Geach objection.

So far, we've been working with the fairly blunt Humean distinction between cognitive and conative mental states. Contemporary quasi-realists such as Gibbard and Blackburn pursue a more subtle account; so it is worth mentioning why I don't think it helps against the present objection.

As we saw briefly in the previous chapter, the first move in the quasi-realist's response to the Frege-Geach worry is to adopt a minimalist stance towards truth. It is fairly clear, however, that minimalism alone won't overcome the worry. For, although minimalism wins the quasi-realist the right to speak of truth-conditions (minimalistically interpreted) and sameness of truth-conditions across asserted and unasserted contexts, this sameness (because it is minimalistic) is nonexplanatory.⁵⁴ Thus, quasi-realists need an account of the type of mental state

⁵⁴ This is made vivid by Dreier (1996), who pointed out that, on a minimalist interpretation of truth conditions, we can define arbitrary predicates usable in declarative sentences to do something other than express a factual judgment (his example is 'is hiyo' which, by stipulation, we can predicate of people to accost them). Sentences deploying these arbitrary predicates will be minimally truth-evaluable. Yet that goes no way towards explaining what is going on when we embed, e.g., "Piers is hyo" in the conditional "If Piers is hiyo, then Charlotte is cute." And the semantic similarity between "Piers is hiyo" in both of these

conventionally expressed by ethical claims that is rich enough both to explain semantic similarity across asserted and unasserted contexts and to mark out an epistemologically and ontologically relevant distinction between ethical and descriptive claims.

So the second move in their response to the Frege-Geach worry is to concede that the traditional conception of a conative mental state won't do the necessary work, but then to appeal to a functionalist philosophy of mind to identify a kind of mental state that can do the necessary work.

The idea is to argue that, whatever we call it, there is a functionalist way to identify a special type of mental state caught up in the psychology of moral motivation; and it is this state that is expressed by our ethical claims. Gibbard (1990)⁵⁵ offers the most worked out account in his characterization of the kind of mental state he calls "accepting a norm":

What is it to accept norms? Acceptance is a state of mind, and a good way to identify this state of mind might be to exhibit its place in a rough psychological theory. The capacity to accept norms I portray as a human biological adaptation; accepting norms figures in a peculiarly human system of motivation and control that depends on language.(Ibid.: 7)

The state of accepting a norm, in short, is identified by its place in a syndrome of tendencies towards action and avowal—a syndrome produced by the language-infused system of coordination peculiar to human beings.(Ibid.: 75)

His idea is that such a system has distinct adaptive advantages and so we should expect it to have evolved wherever people live in communities and have language. He writes,

One way that language influences action and emotions is by letting people think together on absent situations...That opens great new scope for coordination, and so a capacity for shared evaluation would be fitness-enhancing in a species with a complex social life.(Ibid.: 72)

occurrences is as much semantic similarity as minimalism is able to win for the expressivist treatment of ethical sentences in asserted and unasserted contexts.

⁵⁵ To be completely precise, Gibbard (1990) does not yet defend quasi-realism because he doesn't there endorse minimalism about truth. However, in Gibbard (2003) he does endorse minimalism and quasi-realism yet maintains the (rough) correctness of his functionalist account of the noncognitive state conveyed by ethical claims. For more on Gibbard's transition to quasi-realism, see my (2005a).

Insofar as such a system requires the existence of the particular kind of state, there is good evolutionary reason to think that we have such states. Moreover, according to his picture of how this coordination is achieved, there has to be what he calls “normative discussion”, which is the linguistic exchange whereby people move towards consensus about how to think, feel, and act. So the argument is roughly that there are good psychological and evolutionary reasons to think both that there is a special kind of mental state—“accepting a norm”—that figures into a psychology of moral motivation, and that there are linguistic conventions governing a special kind of discourse—“normative discussion”—according to which one expresses the acceptance of norms by producing a normative avowal.

Gibbard’s account of the mental state expressed_(p) by ethical claims is by far the most sophisticated and subtle in the literature. However, even if we grant that ethical claims express_(p) a state of norm-acceptance, the pertinent question is whether this distinguishes them from descriptive claims in the way necessary for the expressivist story. There are three reasons to think that Gibbard’s sophisticated functionalist account doesn’t provide an adequate answer to this question. First, it isn’t clear that this story offers any better response to the Frege-Geach worry. For even if we grant that asserted ethical claims express states of norm-acceptance, it is unclear why should we grant that they express the same mental state in unasserted contexts.⁵⁶ Second, as far as the story goes, a descriptivist could agree that ethical claims express the acceptance of norms but argue that this specially defined kind of mental state is just a special kind of belief. Third, on the present understanding of expression_(p), it seems that there are some

⁵⁶ Gibbard would disagree. In his (1990: 92-102), he develops a model-theoretic semantics for normative claims which he claims is sufficient to give a systematic account of how the state of mind expressed by complex claims deploying normative concepts depends on the states of norm acceptance that would be expressed by the components alone. This account, however, does not overcome the objection made vivid by Dreier’s example (see note 18 above). Gibbard admits this in his (2003: ch 4), and there he approaches the issue differently by developing an general account of what it is for two mental states to stand in disagreement. However, without getting into the details, as far as I can tell, this does not overcome the worry that in embedded contexts, normative claims express a different state of mind from that that they express in nonembedded contexts; so any account of their meaning which rests solely on what state of mind they express will suffer the Frege-Geach problem. See also Dreier (forthcoming).

descriptive claims that can just as easily convey the acceptance of norms as an ethical claim. For example, let's grant that an utterance of "Killing is wrong" conventionally conveys one's acceptance of a norm which forbids killing; wouldn't an utterance of "I accept norms which forbid killing" then also conventionally convey this mental state? But this is (at least sometimes⁵⁷) a descriptive statement about one's own mental states.

The upshot of these reflections appears to be that even if the functionalist account contemporary quasi-realists use to identify the mental state conventionally expressed by ethical claims is correct, this doesn't provide for the sort of expressive contrast between ethical and descriptive claims needed for a successful version of expressivism. In a way, this is a restatement of the challenge of creeping minimalism and valid reasoning that we saw facing quasi-realism in the previous section. Descriptivists will see this as a boon for their argument—even if everything that the quasi-realist says about what is expressed by ethical claims is correct it is still compatible with descriptivism. But I think the problem with quasi-realism really results from a misconstrual of the expression_(p) relation.

The causal and conventional ways of understanding expression_(p) construe it as a relation between utterances and mental states. In order to see that there is an alternative way to understand expression_(p), notice that we sometimes also talk of people expressing commitments by performing particular expressive actions. For instance, in shaking hands, parties to some negotiation can express their commitment to the agreement they have reached; or in signing the confession a murderer can express her commitment to its truth.⁵⁸ The term 'commitment' is

⁵⁷ It might be plausibly questioned whether all self-ascriptions of a mental state are descriptive. For example, a roughly Wittgensteinian account of a typical utterance of "I'm in pain" might construe it as not stating a fact but expressing pain. However, it seems highly implausible to be a Wittgensteinian expressivist about all utterances of this (or any other self-ascriptive) sentence; for, in the context of self-reflection or therapy, these seem to be genuine statements of facts about one's mental states. But even in these contexts, they would conventionally convey the state ascribed.

⁵⁸ In other work in progress, I suggest that this is a third expression relation distinct from the relation between a sentence and a proposition and the relation between the utterance of a sentence and what is conveyed. I argue that it is a relation between a *person* and the *commitment* she expresses by performing

used in several different ways in contemporary philosophy, and, on one of them, it is simply a stand-in for whatever mental state is expressed by a sincerely and seriously uttered declarative sentence. As I mean it, however, a commitment is explicitly and importantly not identical to any particular mental state. For example, we ordinarily think that someone who has made a promise is committed to doing what he has promised to do, regardless of what he wants or believes (or whatever other features of his psychology). All that matters, modulo expression of the commitment, is that he did what it takes to undertake the commitment.

Promises aren't the only speech acts that express commitments; indeed, promises are special because the commitments that come with promises are undertaken voluntarily. Moreover, the normal way that one undertakes the commitments a promise carries is by saying or doing something that simultaneously expresses the commitments undertaken. But not all commitments are voluntarily undertaken, and not all expressions of a commitment are simultaneously undertakings of the very commitment expressed. Someone who asserts "If abortion is not permissible then neither is the death penalty" and "Abortion is not permissible" is committed to the death penalty's not being permissible (or, perhaps, to revising one of her prior assertions). In this case, the speaker is committed regardless of whether she chose or wants to be committed. So, not all commitments are voluntary. Moreover, one can express a commitment that one has already undertaken, as, for example, when one is asked, "Why are you writing that big check?" and one utters the past-tensed: "I promised to pay back the loan by today." In this case, one seems to be expressing a commitment that one already undertook in the past. So, not all expressions of commitments are simultaneously undertakings of the commitment expressed.

an expressive act (such as uttering a sentence). I don't think much hangs on whether this is considered to be a third expression relation or a third way to understand the expression_(p) relation. Either way, what is important is that it provides a level at which a new sort of expressive contrast can be articulated.

There are perhaps several different ways that we express several different kinds of commitments, but notice that one way to pursue a general account of specifically linguistic meaning is by reflecting on the conceptually articulated commitment expressed by certain kinds of linguo-expressive vehicles. This, in turn, provides for a new understanding of $\text{expression}_{(p)}$: $\text{expression}_{(p)}$ can be thought of as a relation between an expressive vehicle with conceptual content (a claim) and the commitment this vehicle articulates or makes explicit.⁵⁹ Roughly, the idea is that in making a claim, one is expressing one's commitment to the correctness of what is claimed. For instance, in claiming that the butler did it, one is expressing a commitment to the butler's having done it. This is almost a tautology, but there is a substantive question about what exactly this commitment comes to. And I think we will make significant progress in understanding the conceptual content of claims by identifying the nature of the conceptually articulated commitments they express.

But how exactly are we to determine the nature of these commitments crucial to understanding the content of our claims? This is interestingly similar to the question: how exactly are we to determine the nature of the mental states crucial to understanding the content of our claims? As we saw, the most promising route for noncognitivism is the quasi-realist pursuit of a functional-role answer to the latter question. That is to say that they try to articulate the nature of the mental state crucial to understanding the content of our ethical claims in terms of its characteristic functional-role in the psychology of moral motivation. But the fact that commitments aren't mental states indicates that such a psychologically specified functional-role

⁵⁹ Brandom (1994) pursues a similar (and much more worked out) account of conceptual content in terms of inferentially articulated commitments. My account differs from his in that it turns on the notion of using claims to *express* commitments, which is different from but related to the notion of *undertaking* a commitment. For example, he suggests that "[s]aying or thinking that things are thus-and-so is undertaking a distinctive kind of inferentially articulated commitment: putting it forward as a fit premise for further inference, that is authorizing its use as such a premise, and undertaking responsibility to entitle oneself to that commitment, to vindicate one's authority, under suitable circumstances, paradigmatically by exhibiting it as the conclusion of an inference from other such commitments to which one is or can become entitled." (2000: 11) I think this may be right, but I doubt that the class of discursive commitments Brandom thinks are undertaken in making a claim is completely coextensive with the class of commitments expressed by one making the claim.

answer to the correlated question about the nature of the commitments would be wrongheaded. Instead, I propose to pursue an inferentially specified *normative-role* answer to the question.

What is a normative-role? A good example of a normative-role comes from chess. Consider what a pawn is. In one sense, a pawn is a physical item in a chess set. However, what makes a particular physical item a pawn and not, e.g., a bishop or a mere piece of marble, is that it is a physical item with a particular normative-role. Specifically, it is the thing (of whatever constitution) whose use is constitutively governed by the rules of chess pertaining to pawns.⁶⁰ Almost anything can, in principal, be a pawn, as long as its use is governed by certain rules. So, these rules are determinative of what a pawn is.

Issues are more complicated when it comes to determining the nature of the relevant sort of commitments in terms of their normative roles. Here, it is not as if we have a rule book for determining the constitutive rules governing different sorts of commitments. However, speaking of judgments, Frege made a helpful suggestion. He claimed that the conceptually relevant way “in which the content of two judgments may differ” is that “it may not...be the case that all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments”(1879: sec. 3). I take it Frege’s idea is that if two judgments differ in their inferential potential, then they differ in conceptual content, whereas, if they only differ in other respects but have the same inferential potential, then they have the same conceptual content.

I want to extend Frege’s idea to the specific kind of commitment expressed_(p) by claims that I think is helpful for understanding the claims’ conceptual content. Thus, the idea is that these commitments are constitutively governed by rules of good inference, which then are determinative of the commitments’ normative-role. As a first approximation, think of the commitment one expresses_(p) by making some claim as a commitment to all of the things that,

⁶⁰ Compare Rosenberg (1974: 39).

according to the logical-conceptual rules of the language in which it is made, materially follow from it and to at least one of the things it materially follows from. For example, if one claims, “Dogwood leaves turn brown in the fall,” then, on this model, one has expressed one’s commitment to the correctness of thinking of some particular tree that if it is a dogwood then its leaves will turn brown in the fall, and of thinking that there is some materially sufficient reason for thinking that dogwood leaves turn brown in the fall.

Of course, many different things might materially follow from any claim. And once we add variable sets of auxiliary premises, we start to get a complex array of potential inferences beginning from any possible claim. Moreover, any claim can be seen as a node in an array of inferences leading to it and inferences leading from it. Ultimately, it is this complex array that individuates the commitments relevant to understanding the conceptual content of claims.

So, the generic idea for this content-relevant sort of commitment is that a commitment to x is something, of whatever constitution, whose use is constitutively governed by rules of good inference pertaining to commitments to x . More specifically, these rules of good inference will determine the inferential potential of the commitment. We can view inferential potential as comprising two components—upstream and downstream potentials. A commitment’s inferential potential is partially determined by a function on all of the combinations of commitments that could count as sufficient reasons for it—this is its upstream inferential potential. And its inferential potential is also partially determined by a function on all of the combinations of things it can count as reason for—this is its downstream inferential potential. (The most natural way, in my opinion, to think about the end of a commitment’s upstream and downstream inferential potentials is in terms of observation and intention—basically, what Sellars (1967) calls “language-entries” and “language-exits”.)

We now have some meat on the bones of the emerging alternative account of $\text{expression}_{(p)}$. Let me summarize. Quasi-realists want to understand $\text{expression}_{(p)}$ in terms of the mental state

conventionally conveyed by a claim. By contrast, the inferentialist understanding of content-relevant commitments encourages thinking of claims as expressing_(p) commitments by virtue of their showing or making explicit what one who makes them is committed to.

The pertinent question to consider at this point is what, on this general model of the inferentially articulated commitment expressed_(p) by statements, we should say about the specific case of commitments expressed_(p) by ethical claims. I think the intuitive motivation behind the sorts of expressivist contrasts pursued by emotivist and quasi-realist versions of expressivism is legitimate, even if these theses views don't stand up under scrutiny. This motivation was to capture the intuitively strong connection between ethical discourse and action in a way that doesn't commit to the existence of strange facts with an objective prescriptive force. And I think the third account of expression_(p) just outlined offers new resources for better capturing this connection in an expressive contrast.

To see this, consider first the distinction between a commitment to a belief and a commitment to an action. For example, in divvying up chores before a party, one person might say, "We still need four bags of ice," and another might say, "I'm on grocery-duty, so I'll get the ice". The first person has expressed what we might think of as a doxastic commitment, while the second has expressed what we might think of as a practical commitment. The former is a commitment to a *factum*—i.e. to something's being true. The latter is a commitment to a *faciendum*—i.e. to acting in such a way so as to make something true.

Now, it's important to recognize there is a way to use the sentence "I'll get the ice" to express a *doxastic* commitment. It would be a commitment to a prediction, on par with "I'll eventually die." However, this is not how we normally use the sentence "I'll get the ice". We aren't typically making a prediction; rather we are, so to speak, making a plan. So, as an initial example of

practical commitments, we should consider this plan-making use of sentences like “I’ll get the ice.”⁶¹

Can this distinction between expressions of doxastic and practical commitments be mapped fruitfully on to the idea that the content of claims is, in general, to be understood in terms of the inferentially articulated commitments that they express? At first, it might not be obvious that it can. For it is common to assume that it is only doxastic commitments that stand in genuine inferential relations, which would mean that it is only doxastic commitments that have inferential potential. It’s common to assume this because there is a neat and natural story to be told about what counts as a good inference between doxastic commitments. The idea is that beliefs aim at the truth, and inference is, at base, a transition between beliefs; so good inference will be the sorts of transitions that are guaranteed to preserve truth.

I’m not sure if this is the correct or complete story to be told about the inferential connections between doxastic commitments, but what is important here is that we clearly need a broader notion of inference if we are to make any sense of the idea of practical reasoning. One way to start to provide such an account is to conceive of inference more broadly as a relation between justifiers and justifieds. In the realm of doxastic commitment, commitment to the belief that there will be 50 people at the party and commitment to the belief that four bags of ice are necessary for a party of 50 people *justifies* a commitment to the belief that we need four bags of ice. Similarly, in the realm of practical commitment, commitment to going to the store to get what we need for the party and commitment to the belief that we need four bags of ice *justifies* a commitment to getting four bags of ice at the store.⁶²

⁶¹ Some writers regiment English here to make the distinction—reserving ‘shall’ for plans and ‘will’ for predictions. See Sellars (1956, 1963) and Rosenberg (2002: 85-86). In what follows I shall adopt this convention.

⁶² This then provides a neat way to frame the question of whether it is preservation of truth that is determinative of the goodness of inferential connections between doxastic commitments. The question becomes: Is the way in which one doxastic commitment justifies another explicable solely in terms of the

It would be nice to have a general account of what makes for good inference in the realm of both doxastic and practical commitments, but, without getting into the vexed issue of what such an account would look like, I think it is clear that there are good practical inferences—conceived of, at least, as a relation between justifiers and justifieds. So, in thinking about the normative-roles of the commitments expressed by various sorts of claims, I think we should recognize both theoretical and practical inference as relevant to inferential potential.

Intuitively, ethical claims such as “One ought not lie” are distinctively practical (at least) in that they are about how we should or shouldn’t act. A plausible way to capture this intuition in the framework of $\text{expression}_{(p)}$ that I have been sketching is by saying that the inferential potential of the commitments expressed by ethical claims has some distinctively practical element. More specifically, I want to suggest that an ethical commitment can serve as the sole practical commitment in a practical inference. For instance, the commitment expressed by “One ought not lie” can be combined with a doxastic commitment such as that expressed by “Falsifying one’s tax forms is a form of lying” to justify a practical commitment to filing true tax forms. Another way to put this point is that a sufficient and legitimate answer to the question, “Why did you file true tax forms?” is: “Not doing so would be lying, and one ought not lie.” We can of course disagree with the ethical or descriptive claim, but if we don’t dispute these, it would normally seem misplaced to continue and ask, “Yeah, ok, but why did you do what you ought to?”⁶³ I think this reveals that the commitments expressed by ethical claims, such as “One

former indicating the truth of the latter. I think there is a genuine debate to be had here between so-called “truth-consequentialists”, such as BonJour (1978), Foley (1987), Alston (1989), and Sartwell (1992) and their opponents, such as Rosenberg (2002) and Kelly (2003).

⁶³ It’s not that this question is devoid of application but that it is not an appropriate question in the normal case of asking for local justifications of actions. As with the two uses of sentences such as “I’ll get the ice,” there is also an explanatory use of these why-questions, where “Why did you do what you ought to do?” is on par with “Why did you get sick?” in that it is used to ask for a third-personal explanatory stance on what happened. (Compare the discussion in Moran (2001: ch. 4) justificatory versus explanatory why-questions.)

ought not lie,” have the downstream inferential potential that is characteristic of practical commitments.

Already, this marks a significant expressive contrast with uncontroversially descriptive claims. With descriptive claims, if someone performs some action ϕ , and we ask her, “Why did you ϕ ?” and she responds with some descriptive claim such as, “Because roses are red,” it doesn’t necessarily seem misplaced for us to continue and ask, “Yeah, ok, but why does that matter?” Of course, sometimes the answer is obvious, as in “Why did you buy the blue car instead of the red car? Because it was cheaper.” But in these cases, unlike the moral case, it seems to me that the context is supplying the implicit further practical reason with which the descriptive claim then combines to provide a reason for action. Perhaps this is one way of capturing the intuition behind Moore’s (1903) open-question argument⁶⁴: it is part of the inferential potential of ethical and descriptive claims that no matter how many descriptive claims you make, you still need some practical claim in order to justify an action, and an ethical claim will do the trick.

So far, I’ve argued that paradigmatically ethical claims such as “One ought not lie” express an inferentially articulable commitment whose downstream inferential potential exhibits important similarities to intuitively practical commitments such as “I’ll get the ice.” Basically, they both can be the action-guiding commitment in a piece of practical reasoning. This downstream difference may be enough of an expressive contrast to capture the sort of intuitively strong connection between ethical discourse and action which motivates noncognitivists. But it seems to me that there is also an upstream contrast implicit in the common view that there is a logical gap between ‘is’ and ‘ought’. Following Hume, many philosophers have thought that one can not infer some ought-claim from any collection of is-claims. There is controversy about the scope of Hume’s principle, but it seems to apply at least to moral ought-claims. Many inspired

⁶⁴ An argument that, of course, I don’t mean to endorse in its original form. See section II of chapter 1 above for my discussion of Moore’s open-question argument.

by Hume take this to be a problem for the view that ethical commitments can be genuinely supported by reasons. However, with the distinction between theoretical and practical inference and a view of ethical claims as expressing practical commitments, we can see that Hume's principle only indicates that ethical commitments cannot be genuinely supported by purely doxastic commitments. Our model of $\text{expression}_{(p)}$ provides a nice way to capture this point. We can say that the upstream inferential potential of an ethical claim must include a higher-order ethical (or at least practical) commitment. This respects Hume's principle without giving up on the thought that many ethical commitments can, in principle, be supported by reasons.

So, the general picture I want to put forward of the contrast between ethical and descriptive statements is this: Both ethical and uncontroversially descriptive statements are ways one $\text{expresses}_{(p)}$ a kind of commitment relevant to understanding the conceptual content of the claim. In both cases, these commitments can be identified in terms of their normative-role in reasoning, which provides for an inferential articulation of the conceptual content of the statements. Both the commitments expressed by making ethical statements and the commitments expressed by making uncontroversially descriptive statements hang together in complex inferential relations. However, there is a contrast between theoretical and practical inferential potentials. Looking upstream, the commitments expressed by descriptive statements don't require anything practical for their inferential support; and, looking downstream, these commitments cannot serve as the practical premise in a piece of practical reasoning. By contrast, when we look upstream in the inferential potential of the commitments expressed by ethical statements, it seems that there must be some practical support; and, looking downstream, an ethical statement can serve as the practical premise in a piece of practical reasoning leading to action.

As we saw, traditional noncognitivists and quasi-realist noncognitivists were both motivated by the distinctive practical nature of ethical discourse. But, without adopting any of their core

theses—(i), (ii), and (iii) (see page 105 above)—I think we have just seen an expressive contrast, articulated in terms of inferential potential, that captures the distinctive practical nature of ethical claims. Ethical claims express_(p) commitments whose inferential potential is distinctively practical. Because descriptive claims do not express distinctively practical commitments in this sense, I'll say that they express purely theoretical commitments, which provides for a new expressivist thesis:

(iv) Ethical claims express_(p) distinctively practical commitments while descriptive claims express_(p) purely theoretical commitments.

Since the nature of these commitments is determined by their inferential potentials, I want to treat (iv) as the foundation of a metaethical position I shall call *inferentialist expressivism*. As I have stressed throughout, not just any expressive contrast can serve as the foundation of a version of metaethical expressivism. We need an ontologically and epistemologically significant expressive contrast. So the question to consider at this point is whether (iv) can provide this.

I think it can. The key is to notice that theoretical and practical commitments exhibit different directions of fit with the world. What does it mean to say that they exhibit different “directions of fit with the world”? In the context of expressing commitments, I understand this metaphor in a related although importantly different way from most who use it. Most philosophers working with this metaphor use it to articulate the psychological distinction between beliefs and desires, which is a crucial prerequisite to answering the question in the theory of motivation of whether, as Humeans would have it, motivation always also requires a desire or, as anti-Humeans would have it, a belief alone can motivate action. For example, in the course of arguing against Humeanism, Platts writes:

The distinction is in terms of the direction of fit of mental states to the world. Beliefs aim at being true, and their being true is their fitting the world; falsity is a decisive failing in a belief, and false beliefs should be discarded; beliefs should be changed to fit with the world, not vice versa. Desires aim at realization, and their realization is the world fitting them; the fact that the indicative content of a desire is not realized in the world is not yet

a failing in the desire, and not yet any reason to discard the desire; the world, crudely, should be changed to fit with our desires.(1979: 257)

While, in the course of arguing for Humeanism, Smith writes,

[T]he difference between beliefs and desires in terms of their direction of fit comes down to the difference between the counterfactual dependence of a belief and a desire that p, on a perception that not-p: roughly, a belief that p is a state that tends to go [out] of existence in the presence of a perception that not-p, where a desire that p is a state that tends to endure, disposing the subject in that state to bring it about that p. Thus, we may say, attributions of beliefs and desires require that different kinds of counterfactuals are true of the subjects to whom they are attributed. We may say that this is what difference in their direction of fit is.(1987: 54)

A reason for appealing to (and then trying to cash out) the metaphor of direction of fit, in both cases, is that the ordinary notion of ‘desire’ is too narrow to capture what is at issue between Humeans and anti-Humeans. Intentions, plans, or even mere inclinations may all serve as “pro-attitudes” whose combination with a belief motivate action, but this shouldn’t be thought to undermine Humeanism straightaway. So proponents and opponents of Humeanism try to come up with a neutral characterization of ‘desire’ (broadly construed) in order to frame the question of whether a desire (broadly construed) is always necessary for motivation.

The most noteworthy difference between Platts’ and Smith’s characterizations of the distinction between direction of fit is that Platt’s is stated in normative-epistemological terms while Smith’s is stated in causal-psychological terms. For the project of codifying the folk-psychological practice of explaining action into a theory of motivation, I suspect that Smith’s strategy is more promising.⁶⁵ And this project is intimately tied up with the noncognitivist attempt to treat ethical claims as expressive of nonbelief desire-like mental states, in order to

⁶⁵ Although the viability of his specific characterization of directions of fit has, I believe, been undermined from two different directions. Schuler (1991) argues that there are some propositional attitudes—e.g. hope that p—which can combine with a belief to motivate action but which meet Smith’s characterization of a belief-like direction of fit. Humberstone (1992) argues that Smith’s use of the notion of ‘perception’ can be read either doxastically or nondoxastically. If the former, then Smith must presuppose a form of the very distinction he is trying to explain. If the latter, then Smith hasn’t truly captured the distinction between beliefs and desires broadly construed since there are some beliefs that p that don’t tend to go out of existence in the face of a perception that not-p—i.e. those we hold when we don’t take our perceptions at face value.

underwrite internalist intuitions while simultaneously gaining traction against Mackie's argument from queerness. However, here, I am no longer specifically interested in the theory of motivation or the psychology behind ethical discourse, but instead in the various types of inferential networks in which different types of claims are caught up. Thus, Platt's normative-epistemological characterization of the difference in direction of fit between beliefs and desires strikes me as a better place to begin to understand the difference in direction of fit between theoretical and practical reasoning.

Importantly, Platt's normative characterization of this distinction draws on Anscombe's original use of the metaphor to characterize the difference between a man's shopping list and a private detective's list of items bought by the man:

...if the list and the things that the man actually buys do not agree, and this alone constitutes a mistake, then the mistake is not in the list but in the man's performance...whereas if the detective's record and what the man actually buys do not agree, then the mistake is in the record.(1957: 56)⁶⁶

Here, however, we should note that the idea of a difference in direction of fit doesn't attach to two different kinds of propositional attitudes—beliefs and desires (broadly construed)—but to two different kinds of lists.⁶⁷ Moreover, in contrast to what Platt's metaphorical talk of the potential “failings” in beliefs and desires to achieve their “aims” might suggest, here the

⁶⁶ It appears that the *original* use of the metaphor was by Austin (1953) to characterize two different kinds of speech-act involving predication. For instance, “The car is blue” and “The blue car is to be washed first.” Humberstone (1992) argues that this is a completely unrelated distinction to the one which Anscombe, Platts, and Smith try to draw, since it differentiates two ways of using what might be thought of as sentences aimed to fit the world, rather than two ways propositional attitudes may be related to the world. However, I'm not so sure that Austin's use of the distinction is so unrelated to Anscombe's use in the present example. For notice that we could think of the two lists as both comprised of predicative sentences: “Apples are to be bought; milk is to be bought...” and “Apples were bought; milk was bought...” But that doesn't mean that they don't exhibit different directions of fit. Since, in general, I am less interested in the psychological question about how to differentiate between different kinds of propositional attitudes and more interested in the normative-pragmatic question of how to differentiate between different kinds of reasoning and their articulation in expressive behavior, I see Anscombe's list example and Austin's speech-act example as both similar and more illuminating.

⁶⁷ To be sure, Anscombe ultimately wants to press this distinction into an account of the difference between intentions and beliefs, which she does see as types of propositional attitudes. And while it may be made to work in this specific case, as Velleman (1992/2000) has argued, it cannot in general distinguish cognitions from conations because there are many plausible “no fault” failures of fit going both ways.

evaluative notion of a “mistake” attaches to more recognizably genuine bearers of responsibility—people engaging in specific teleologically structured performances.

I think we can view both kinds of lists as instances of expressions of commitments. The man’s shopping list can be seen as the expression of a practical commitment to engage in a particular behavior—viz. to buy particular items at the grocery—while the detective’s list can be seen as the expression of a theoretical commitment to the obtaining of a particular state of affairs—viz. that the man bought particular items at the grocery. We can gain a better understanding of the difference between theoretical and practical commitments by reflecting on the difference in the kinds of mistake involved in failing to meet a commitment and, correlatively, the different kinds of responsibility that attach to each kind of commitment.

Sticking with Anscombe’s example and assuming that both lists are indeed mistaken, what is the difference in the kind of mistake? One way we can characterize these errors is in terms of what would have to be done to correct the mistake. For example, if the lists both include “celery” but the shopper bought no celery, then the detective can correct his list by erasing the line that says “celery” and this should be enough to satisfy his employer. However, clearly it won’t do for the man arriving home without the celery necessary for tonight’s dinner to erase the line that says “celery”—his wife will probably expect him to go back to the store. More generally, it seems that a theoretical commitment can, so to speak, itself be mistaken, in which case one who makes the commitment can correct this mistake by changing the content of the commitment; however, practical commitments aren’t themselves, in this way, the mistake, thus correction requires performing some other act than changing the content of the commitment.

In my view, this difference in direction of fit between theoretical and practical commitments is, so to speak, the reflection of upstream and downstream differences in the inferential articulation of the two types of commitments. Claims expressing theoretical commitments require theoretical justification and obligate/license theoretical implications, while claims

expressing practical commitments require practical justification and obligate/license practical implications. And if we adopt the inferentialist idea that the meaning of a claim can be articulated (at least in part) in terms of its normative-role in reasoning, we have the beginnings of a story about a genuine semantic difference between theoretical and practical claims.

But what about “mixed” claims such as, “Either he has a media pass or he shouldn’t be on the court”—is this a theoretical or a practical claim? It seems like both, but if we say it is both, then it may seem that the distinction between claims which express theoretical commitments and claims which express practical commitments does not really represent a difference in direction of fit with the world.

This is a reasonable but answerable worry about the account I have been sketching. The best response to it would be to move the discussion from the level of full claims to the level of the component concepts which comprise full claims, and then to say that the theoretical or practical nature of a concept is a matter of the types of claims it can be used to make and their inferential roles. Thus, the claim “Either he has a media pass or he shouldn’t be admitted” is both theoretical and practical because it mobilizes both theoretical and practical concepts; but we can break the claim down into various parts whose differential connection to theoretical and practical inference is clear.

However, even if one grants, that ethical claims express practical commitments rather than theoretical commitments, presumably not all expressions of practical commitments are ethical claims. So there remains a question about how to distinguish the ethical within the realm of the practical. This is a difficult and large question. In order to put a little more meat on the bones of inferentialist expressivism, I want speculate about how an inferentialist might further articulate the precise kind of practical commitment expressed by ethical claims. Although I don’t see the line I take as strictly necessary for being an inferentialist expressivist, I think it will add

plausibility to the account of ethical claims developed thus far and aid in meeting a number of potential objections.⁶⁸

Even if it is necessary that an utterance express a practical commitment for it to be an ethical claim, this condition surely isn't sufficient—for there are plenty of practical commitments that are not ethical. For example, although they seem clearly practical, instances of the claim, “We should put the cake in the oven” are probably very rarely expressions of specifically ethical commitments. Is there a way to isolate those practical commitments that are specifically ethical? I doubt that there is a way to do this precisely without moving from the arena of metaethical theorizing to the arena of normative ethical theorizing. This is because proposed accounts of which class of practical commitments are specifically ethical are likely to have specific normative ethical consequences or draw on specific normative ethical presuppositions. Nevertheless, in this section, I want to sketch an account of the specifically ethical commitments in a way that is relatively but not completely neutral on normative ethical issues.

There is much debate about the nature of practical reasoning and this debate clearly bears on my proposed distinction between ethical and descriptive claims. One debatable, but I think correct principle about practical reasoning is that it, unlike theoretical reasoning must at least implicitly involve a practical premise in order to validly reach a conclusion about what to do. To see this consider two pieces of reasoning—one theoretical and one practical: Sometimes I want to know whether my car needs a particular repair, so I take it to the shop and my mechanic explains that unless I replace the timing belt, the car will soon break down. I might double-check what he says with some of my friends who know more about cars than I do. If they agree with

⁶⁸ One could see what follows as the parallel to Gibbard's (1990) attempt to go beyond the mere noncognitivist thesis that ethical claims express nonbelief attitudes by arguing that there is a specific type of nonbelief attitude—norm-acceptance—which ethical claims express. Of course, he pursues this via a speculative evolutionary psychology. I shall pursue my analogue of this via a speculative normative pragmatics.

the mechanic, I conclude that my car will break down unless I replace the timing belt. This is theoretical reasoning leading to a theoretical conclusion about what is the case. However, I might also think that replacing the timing belt is relatively inexpensive and that I need a car to get around, so decide to replace the timing belt. This is practical reasoning leading to a practical conclusion about what to do. If we set out the pieces of reasoning, we get something like:

Theoretical

- (T1) My mechanic and others agree that unless I replace my timing belt, my car will soon break down.
- (T2) If my mechanic and others agree that unless I replace my timing belt, my car will soon break down, then unless I replace my timing belt, my car will soon break down.
- (TC) Thus, unless I replace my timing belt, my car will soon break down.

Practical

- (P1) I need my car to get around.
- (P2) Unless I replace the timing belt my car, which is my necessary means for getting around will soon breakdown,
- (PC) Thus, I shall replace the timing belt.

Now, the theoretical piece of reasoning seems straightforwardly valid. If the premises are true, then there can be no question as to whether the conclusion is true; however, the practical piece of reasoning seems to rely on an implicit premise for its validity. It needs something such as,

- (Po) I shall do whatever is necessary to get around,

which makes explicit the practical premise necessary to reach the practical conclusion at (PC) validly. So, with regard to the question of whether a particular practical commitment, e.g. that

expressed by (PC), expresses an ethical commitment, one place to look for an answer is at the practical premises, e.g. (Po), on which it implicitly rests. If it rests on an ethical premise, then we might say that, *ceteris paribus*, it is also an ethical commitment. But this just pushes the problem back. How do we tell whether a practical premise is ethical?

The two most prominent kinds of theories of right action are consequentialism and deontology. Both appeal to the universality of specifically moral reasons to engage in particular actions. Now, in one sense, all reasons are universal: my reason for, e.g., going to the store is to get bread, but anyone who wanted bread and found herself in my situation would have a reason to go to the store. It is, so to speak, in the nature of reasons to be hypothetically universal. However, consequentialists and deontologists both think moral reasons are universal in a stronger sense: they are nonhypothetically universal, or as it is often put, categorical. This means that they are reasons for someone to behave in a certain way, irrespective of her antecedent desires. For example, both the consequentialist and the deontologist will agree that there may be specific practical reasons for me not to steal on some specific occasion—e.g. that the likelihood of my getting caught and incarcerated is high—and these reasons are hypothetically universal. But, for both consequentialists and deontologists, the moral reason I shouldn't steal is, whatever it is, the same for everyone regardless of one's antecedent desires or particular situation. That is, that doing so doesn't maximize utility or requires acting on a maxim that cannot be willed as universal law. These considerations, unlike the fact that the likelihood of my getting caught and incarcerated is high, are moral reasons, if at all, regardless of antecedent desires. So, one appealing way to differentiate moral commitments from other practical commitments is on the basis of the implied categoricity of either the commitment itself or the reasons that one takes to justify the commitment.

In laying out the practical syllogism above, I used the word 'shall' to indicate that the commitments expressed at (Po) and (PC) are practical commitments to courses of action rather

than theoretical commitments to descriptive predictions. Sticking to this regimentation of English (compare note # above), we might distinguish between individual and collective practical commitments by using the first-person singular and the first-person plural, respectively. So, (Po) and (PC) both express individual practical commitments, but it's possible that (Po) itself is based on a more general piece of practical reasoning, e.g.:

(P-3) We shall each provide certain basic goods for our own children,

(P-2) Thus, I shall provide certain basic goods for my children.

(P-1) In order to provide certain basic goods for my children, need to be able to get around,
thus

(Po) I shall do whatever is necessary to get around,

in which case, (P3) is ultimately grounded on a collective practical commitment. Moreover, as such, it starts to seem more like the expression of a specifically ethical commitment.

Of course, a lot will depend on the extension of the 'we'. If it is merely a particular sub-culture, the ethical character becomes questionable. For example, if (P3) were ultimately grounded on the collective commitment expressed by:

(P-3') We upper-class white southerners will provide our children transportation to debutant balls and dancing lessons,

from which follows,

(P-2') Because I am an upper-class white southerner, I shall provide my children transportation to debutant balls and dancing lessons,

which with,

(P-1') In order to provide my children transportation to debutant balls and dancing lessons,
I need to be able to get around,

implies (P₀) and ultimately (P₃), then we might worry that (P₃) isn't ethical after all. But if the 'we' in (P-3) is all rational agents, then the sort of practical premise on which (P₃) is grounded would manifest the universality that both consequentialism and deontology suggest is necessary for a practical reason to be a specifically ethical reason.

If this line of thought is on the right track, then we might say that ethical claims are the expression of practical commitments that either themselves are taken to be collective practical commitments where the collective is all rational agents or are taken to be grounded on such collective practical commitments. In the regimented language of the syllogistic representation of practical reasoning, these collective practical commitments get expressed as instances of

(M) We all shall...

where the 'we all' indicates all rational agents and the 'shall' indicates that it is a practical commitment—one, that is, which can figure in the right way in reasoning to a decision about what to do and, if supportable by further reasoning, this further reasoning will also be practical—rather than a predictive theoretical commitment. Notice that this is not to say that in order for a claim to be ethical it must express a practical commitment that can in fact be justified with something of the form (M); rather the person making the claim has to intend it as so grounded. This is a reflection of the fact that we might not actually have moral reason to act in the way we take ourselves to have moral reason to act.

So, from the imprecise suggestion that ethical claims express practical rather than theoretical commitments, we've reached a point where I can articulate the particular version of inferentialist expressivism I want to propose:

Ethical claims express practical commitments that are intended as of the form (M) or as grounded on commitments of the form (M), while descriptive claims express theoretical commitments.

The idea is not that all ethical claims have a surface grammatical form like (M). Clearly, many ethical claims involve no explicit ‘shall’-operator or use of the first-person plural. Rather the idea is that the inferential articulation and thus semantic content of ethical claims is practical, and the claims that count as specifically ethical are the ones that express a practical commitment of the form (M) or of something intended⁶⁹ to be grounded in a commitment of the form (M). So, for example, a claim such as “Murder is wrong” might, on this analysis, be taken to have the logical form of “We all shall not commit murder”, whereas a claim such as “I ought to visit my ailing mother” might, on this analysis, be taken to have the logical form of “I shall under appropriate circumstances visit my ailing mother” where this is taken to be grounded on some more general practical commitment of the (M) form such as, “We all shall respect and care for our parents.”⁷⁰

One might wonder whether this version of expressivism is a version of cognitivism or noncognitivism. My answer is that it’s neither really. It does articulate an epistemologically and ontologically relevant distinction between what is expressed by ethical claims and descriptive claims, but it doesn’t make this distinction, like traditional noncognitivists, in terms of whether the claims express propositions or not, nor does it make the distinction, like quasi-realists, in terms of whether the claims express cognitive or conative mental states. Rather it makes it in terms of the type of commitment expressed. But as with quasi-realism, this can be no mere pragmatic matter; it has to have semantic import in order to serve in an expressivist answer to Mackie’s argument from queerness. And, with the help of the normative-role semantics sketched above, it can. For I have explained the difference between theoretical and practical commitments in terms of their different directions of fit with the world, which in turn was

⁶⁹ How strong is my talk of “intending” one’s commitments to be grounded in particular ways meant to be taken? I understand this notion fairly weakly, such that its enough if the person “intending” her commitment to be so grounded is disposed in ideal circumstances under appropriate questioning to bottom out with something of the form (M) in the giving of reasons for her original claim.

⁷⁰ This line of thought owes its inspiration though not exact details to Sellars (1968: ch. 7).

explained as a matter of manifesting different normative roles; and, as long as the semantic content of the claims expressing these commitments is partially constituted by these normative roles, we have an account of the crucial expressive difference between ethical and descriptive claims.

Is this account open to the objections I pressed above against the quasi-realist's account of expression_(p) and their correlative attempt to mark out an expressive contrast between ethical and descriptive claims? I don't think so. The first problem raised above was that even some descriptive claims express_(p) conative mental states, so the mere fact that a claim expresses_(p) a conative mental state isn't enough to make it nondescriptive. With the normative construal of expression_(p) in terms of the commitments expressed, I think things break exactly the opposite way. Given that a practical commitment has a different direction of fit with the world from a theoretical commitment, I think it is fair to say that the fact that a claim expresses_(p) a practical commitment is enough to make it nondescriptive. To be more precise, some claims may express *both* practical and theoretical commitments (this is perhaps the most natural way to think about claims involving thick ethical concepts, such as 'courage' and 'chaste'). In this case, the right thing to say is probably that the fact that they express both a practical and a theoretical commitment indicates that they are not *purely* descriptive; and careful philosophical analysis can pull the two parts apart.

The second problem raised for the quasi-realist account of expression_(p) was an updated version of the Frege-Geach problem. Here too, I think my account brings new resources to the table. By locating the expressive contrast at the level of the commitment expressed rather than at the level of the mental state expressed, an inferentialist expressivist doesn't have to commit to there being any particular kind of noncognitive mental state present when one embeds an ethical statement in a conditional. According to inferentialist expressivism, the nature of the commitment expressed by any statement is a function of the combination of things that count as

reasons for it and the combination of things it counts as reasons for. Part of this inferential potential will include the possibility of combining the commitment with a logically more complex commitment that embeds the original commitment. For instance, the inferential potential of the commitment expressed by, “One ought not lie,” will include its possible combination with the statement, “If one ought not lie, then one ought not get one’s little brother to lie,” to support the further statement “One ought not get one’s little brother to lie.” One doesn’t have to believe or feel that any of this is true, in order for the statement to express a commitment that is individuated by a pattern of good reasoning of which this toy inference is a small part. By taking the debate out of the realm of functional-role psychology and into the realm of philosophic reflection on normative-roles, we gain new resources, within an expressivist framework, for theorizing about the logical relations of different sorts of statements.

Because this account of expression_(p) avoids the problems I pressed against quasi-realism above, I am hopeful that it makes real progress in metaethical debate. However, the real test is whether it can help to make progress on the four complex challenges that emerged in our discussion thus far. So it is to this that I now turn.

In the previous chapter I articulated three complex challenges that emerge out of careful reflection on the most prominent descriptivist accounts of the meaning of ethical claims currently on offer in metaethical discussions. And earlier in this chapter, a fourth challenge emerged which attaches to any expressivist attempt to account for the meaning of ethical claims. I want to argue now that the inferentialist version of expressivism that I have just sketched has the resources to meet all four challenges.

Mackie’s challenge came in the form of his argument from epistemological and ontological queerness against ethical truths. As we saw, nonnaturalists reject Mackie’s assumption of naturalism, which is a pretty drastic move unless one can show why it is legitimate to posit

nonnatural facts, which in turn seems to require a plausible epistemology for how we can, in principle, come to know nonnatural facts. But here we found even contemporary defenses of nonnaturalism to be lacking. Shafer-Landau's moral epistemology doesn't clearly explain epistemic access to the nonnatural ethical facts at all. Regan's moral epistemology stresses a plausible point that ethical knowledge is peculiarly practical—it's knowledge what to do—and its possession seems to require of conceiving of ourselves as practical agents. However, it was mysterious why this requires its possession to be epistemically connected to a special realm of nonnatural facts.

This left relativism and naturalism as alternative descriptivist answers to Mackie's challenge, and I suggested that these are although not completely intuitive, nonetheless sufficient answers, if the theories don't generate other problems—which, of course, as we saw and will see again below, they do. Traditional noncognitivists pursue a different answer to Mackie. They deny that ethical claims are truth-evaluable or expressive of cognitive mental states, which opens up space for the argument that ethical claims shouldn't even be construed as describing some independent ethical reality. With this move, they were able to respond to Mackie by arguing that what constitutes the meaning of ethical claims is something such that making an ethical claim doesn't commit one to any sort of corresponding fact in the world. To this end, they seemed forced to adopt a psychologistic approach to semantic explanation, which is why they view the thesis that ethical claims express_(p) conative mental states rather than beliefs as enough to show that ethical claims are no more ontologically committing than commands or mere evincings of emotions. And, under the typical way of understanding the direction of fit metaphor (as a way to distinguish beliefs from desires), we can how this is supposed to work: beliefs are states that are aimed at fitting the world, so expressing a belief is taking a stand on how the world is, but conative mental states are states that are rather aimed at changing the world to fit them, so expressing a conative mental state doesn't involve taking a stand on the how the world is.

Contemporary quasi-realist versions of noncognitivism try to piggy-back on this traditional noncognitivist response to Mackie's challenge but with the twist of admitting the (minimalistically construed) truth-evaluability of ethical claims. This means that their entire answer to Mackie rests on the viability of their semantic interpretation of the idea that ethical claims express a conative state of mind. And this generated the minimalist's challenge, which I'll discuss in more detail below. Moreover, as we saw above, it is difficult to defend the thesis that what differentiates ethical claims from descriptive claims is that they express conative mental states, which is a problem for both varieties of noncognitivism.

With the third normative account of $\text{expression}_{(p)}$ and inferentialist version of expressivism that it allowed me to sketch, the focus is changed from what type of mental state is expressed to what type of commitment is expressed. I've suggested that it's highly plausible to think that ethical claims express a special sort of practical (rather than theoretical) commitment. Moreover, this expressive difference is both semantically relevant and epistemologically and ontologically significant. It's semantically relevant on the normative-role approach to semantic explanation that I've adopted, which construes the semantic content of all claims in terms of the normative roles of the commitments they express. It's epistemologically and ontologically relevant because practical commitments manifest a different direction of fit with the world from theoretical commitments. Interestingly, however, this difference in direction of fit is different from but related to the putative difference in direction of fit manifested by conative mental states and beliefs. In my view, both practical and theoretical commitments have both upstream and downstream inferential implications, so there is a sense in which both are objects of potential justification and are themselves potential justifiers. But the type of justification in play is different, and this difference makes the key difference in epistemological and ontological commitment. A theoretical commitment requires and provides evidence about how the world is, which indicates that it is intended to fit the world; in contrast, a practical commitment requires and provides practical grounds for action, which indicates that one aims to have the world fit it.

Thus, the rough idea is that one who expresses a practical commitment is not, in the first instance, committed to the world being any particular way, but is committed to making it a particular way.⁷¹ However, this is not because what is expressed by practical claims is not the object of potential justification and reasoning (as noncognitivists typically suggest).

Thus, with my inferentialist version of expressivism now on the table, we can see how to preserve the upshot of the noncognitivist's answer to Mackie—viz. that ethical claims are not descriptions of some ethical reality—without endorsing the thesis that ethical claims are noncognitive. Independently of whether ethical claims express cognitive or noncognitive mental states, they might nonetheless express something nondescriptive. Specifically, I've argued that they express a special sort of practical commitment, and practical commitments manifest a different (=nondescriptive) direction of fit with the world from theoretical commitments. If this is right, we don't need to claim that ethical claims describe some extant ethical reality in order to be cognitivists. So, we can have a perfectly general and satisfactory account of the meaning of ethical claims which treats them as legitimate without committing to the existence of queer facts that might raise Mackie's naturalistic qualms in the first place.

One might worry that even if this line of argument avoids viewing ethical claims as ontologically committing, it rides on another ontologically troublesome issue—viz. the ontological status of a commitment. I've cashed out the nature of commitments in terms of their

⁷¹ I say "in the first instance" because I think there is a further sense in which practical commitments *do* carry ontological commitment, but not of the kind that Mackie objects to. This shows up in the psychologist distinction pursued by Bratman (1987) between plans and desires. He argues that intentions are not reducible to beliefs and desires because they manifest a plan-like character that is itself not reducible to beliefs and desires. Plans, unlike desires, are things from which a subject and his community can *inter alia* plan further. So, for example, if I plan to meet Kristin for lunch, I can then take it as a settled fact (=am ontologically committed to the fact) that I won't be in the apartment at 12:15, and thus tell my mother that she shouldn't call back between 12 and 2, and we can plan to talk at some later time. The mere desire to meet Kristin for lunch doesn't carry this ontological commitment to a fact about what I will do in the future—for, Kristin might be out of town, in which case I know that the desire won't be satisfied. Likewise, I think practical commitments *do* ontologically commit one who has them to some (perhaps weakly conditional) fact about future behavior. But, while this distinguishes the inferentialist expressivist's answer to Mackie from the noncognitivist's, it clearly does not represent an ontological commitment carried by ethical claims that Mackie should find troublesome.

inferential roles—i.e. in terms of what can justify them and what they can justify. However, one may reasonably worry that claims about what justifies what are also normative claims and, as such, invite similar ontological worries as ethical claims. There are two things to say about this worry.

First, if my account is successful in trading the ontological worries generated by ethical discourse for ontological worries generated by discourse about commitments and inferential roles, I think that is real progress. For all parties to the debate should admit that there are some commitments—even if there is then debate to be had about what their nature is. For example, I take it everyone will agree that one who signs a contract is, *ceteris paribus*, committed to its terms. Is that a queer fact? Perhaps, but it is less queer than the fact that murder is wrong (at least as this fact is typically understood). We should, of course, admit that there is a difficult philosophical project of spelling out the nature of commitments in naturalistically acceptable terms; I haven't even begun to broach this project. However, I don't see why it should be thought to be impossible. Commitments carve up a social-pragmatic space, and the concepts crucial to understanding the importance of the contours of this space may not be reducible to the descriptive concepts of naturalistic explanation. But this in no way implies that the *stuff* which constitutes this space is nonnatural. In my opinion, we should distinguish between what might be called “methodological naturalism” and “ontological naturalism”. The former says that the only legitimate methods for pursuing understanding are the roughly descriptive-predictive methods of the natural sciences (whatever exactly these happen to be). The latter says that the stuff of our world is stuff which is, in principle, investigatable by the natural sciences (whatever these happen to be). I take Mackie's worry to be at the level of ontological naturalism, and I think inferentialist expressivism avoids this worry insofar as it is possible to see commitments as constituted by (perhaps highly complex) natural stuff. And there is no reason to deny this condition—except, perhaps, perhaps for reasons stemming from a commitment to methodological naturalism, which I think is incorrect.

Second and more significantly, the inferentialist version of *ethical* expressivism that I have sketched here can be seen as a model for explaining away the ontological significance of normative claims more generally. I won't argue this fully here, but if we conceive of claims about what justifies what as expressing practical commitments to certain *inferential* behavior, then, because practical commitments exhibit a different direction of fit from theoretical commitments, we can see those inferential commitments as not ontologically committing.⁷²

The second and perhaps most complex challenge was the *action-guiding challenge*. This was the challenge to be consistent with and preferably to provide the resources to explain the truth of the positions I called rationalism and (weak) internalism. This came out first in conjunction with evaluating contemporary naturalist positions, but it is a thorn in the side of each of the alternative accounts we've considered. We can see its force by noticing that ethical naturalism seems to be motivated by three things. (i) the view that ethical facts are universal in their application; (ii) the assumption that ethical discourse is cognitive in that ethical claims are truth-evaluable and expressive of beliefs, which is taken to imply that ethical claims aim to describe independent ethical reality; and (iii) a commitment to the naturalist thesis that the only real facts are natural facts. However, the strategy of contemporary naturalists for addressing these concerns leaves one wondering in what sense the naturalistically respectable ethical facts are ethical. There are two reasons for this: (iv) As naturalists construe them, ethical facts aren't obviously the sorts of facts that could be intrinsically reason giving, which makes it hard to square naturalism with the rationalist idea that endorsing an ethical claim is ipso facto to take there to be a justifying reason for action. (v) As naturalists construe them, it is doubtful that

⁷² A full dress account would, I think, need to distinguish between hypothetical norms and categorical norms. We can view many normative claims as theoretical claims about what are necessary means for achieving some end. For instance, when one wants to stay healthy, one ought not smoke. But commitment to this normative truth can be seen as a theoretical commitment to the fact that the best way to stay healthy is to avoid smoking. The normative claims that will need special treatment in terms of expressing practical commitments are the ones that cannot be construed as hypothetical. (Thanks to Graham Hubbs for pressing me to spell out this point.)

recognition of ethical facts has any intrinsic motivational force, which makes it hard to square naturalism with the (motivation) internalist idea that ethical judgments are intrinsically motivating. The third challenge is to make a plausible story about (i)-(iii) consistent with a plausible story about (iv)-(v).

There is a lot of room for maneuvering on either side of this debate due to the fact rationalism and internalism admit of various interpretations. However, this much seemed clear from our discussion in chapter II: current forms of nonreductive a posteriori naturalism take a too drastic stance vis-à-vis (iv) and (v) to meet this challenge. We saw that each of the other two descriptivist accounts—relativism and nonnaturalism—have a somewhat more nuanced stance vis-à-vis (iv) and (v), but achieve this by flat out rejecting one of the motivations for naturalism. Relativists reject (i) and nonnaturalists reject (iii). These moves generated independent problems but they also seem to require explaining away pretty strong intuitions that lead to (i) and (iii), which leaves the relativist and nonnaturalist responses to this second challenge pretty suspect. Noncognitivists reject the other motivation for naturalism—viz. (ii). And we saw that by arguing that ethical discourse is noncognitive, they can develop a highly satisfactory story about (iii) which lets them endorse a strong version of (v) internalism; and they at least have some plausible things to say about why we think (i) is true. However, in addition to other problems attaching to noncognitivism, this all comes at the cost of requiring a strongly revisionary stance towards rationalist intuitions. I think this failure to tell a plausible story at (iv) renders the noncognitive answer to this third challenge fairly suspect.

Like noncognitivism, my version of inferentialist expressivism also rejects (ii)—but not completely. Notice that (ii) has two parts—viz. (a) the assumption that ethical discourse is cognitive in that ethical claims are truth-evaluable and expressive of beliefs, and (b) which is taken to imply that ethical claims aim to describe independent ethical reality. These two parts do not have to stand or fall together. Depending on what is meant by “cognitive”, “truth-evaluable”,

and “expressive of beliefs”, there is room to endorse (a) while rejecting (b). Specifically, I’d like to suggest that any discourse in which typical claims are within the purview of reason and the possibility of knowledge is cognitive. And this is consistent with a fairly minimalist conception of truth-evaluability and expressiveness of belief. Given this understanding of (a), one need not reject (a) in rejecting (b). For, as long as we allow that practical reasoning is a form of reasoning and knowing what to do is a form of knowledge, ethical discourse can, on my inferentialist account, be seen to be within the purview of (practical) reasoning and the possibility of (practical) knowledge. Its claims will be truth-evaluable and expressive of beliefs, on the minimalistic interpretations of these terms. But all of this is consistent with insisting that, because ethical claims are expressive of practical commitments and practical commitments have a different direction of fit with the world from theoretical commitments, ethical claims do not aim to describe an independent ethical reality.

In addition to this account of (ii), I think the inferentialist expressivist has a plausible story to tell about (iii); and because of my specific story about ethical claims expressing collective practical commitments, I think inferentialist expressivism has the resources for a satisfying story about (i). Moreover, the view is very compatible with rationalism since to be practically committed just is to think there’s a reason for action. It is also compatible with the most plausible (weak) form internalism that emerged in our discussion of relativism and noncognitivism. So, my inferentialist version of expressivism hedges on (ii) by divorcing (a) and (b). And otherwise it has a pretty plausible story to tell about (i), (iii), (iv), and (v). In my opinion, that’s pretty good and a much better answer to the action-guiding challenge than any of the other accounts have to offer.

The third challenge was the *relativist’s challenge*, which came from reflection on relativist accounts of the meaning of ethical claims. Relativism is clearly motivated by the combination of the apparently robust and recalcitrant diversity of sincerely endorsed ethical claims and the

perceived failure of universalist construals of ethical claims to plausibly explain this lack of consensus. However, the relativists' strategy for explaining this diversity is drastic: they suggest that the diversity doesn't represent genuine disagreement. And while there is room for the relativist to maneuver to capture some genuine disagreement (by widening beyond the individual subject's desires and stances the basis to which ethical claims are held to be relative), it seems that whatever the basis is thought to be, we can construct cases where the relativist has to counterintuitively treat apparently genuine disagreement as mere logically independent difference in opinion. So the relativist's challenge was to come up with an account of the meaning of ethical claims that (i) has the resources to explain the apparently robust and recalcitrant diversity of sincerely endorsed ethical claims while (ii) avoiding the counterintuitive result of treating apparently genuine disagreements as mere logically independent differences of opinion.

We found that both naturalists and nonnaturalists have to punt here and chalk up the apparently robust and recalcitrant diversity of sincerely endorsed ethical claims to some questionable epistemological failure, whose relation to specifically moral epistemology is doubtful. Noncognitivists, on the other hand, have a more plausible way to simultaneously meet (i) and (ii), which stems from their core thesis that ethical claims express conative mental states. The idea was that different people have different moral attitudes, but unlike beliefs these attitudes aren't aimed at tracking some antecedently existing moral reality. Because of this lack of regulation for the independent ethical facts, we shouldn't expect moral attitudes to converge, and this explains the robust and recalcitrant diversity of ethical opinion. However, expressions of conative mental states can nonetheless represent a sort of disagreement—it's a "disagreement in attitudes." And it is the uptake of this sort of disagreement into our expressive behaviors that the noncognitivist uses to capture the sense in which apparently genuinely conflicting ethical claims really do represent a genuine disagreement.

A relativist is likely to be unimpressed with such “disagreement in attitudes”. For it seems to be the wrong sort of disagreement to explain the intuition of disagreement in ethical cases. It’s hard to articulate what exactly is wrong with it, but part of the problem is surely that we take our ethical claims and disagreements to be regulated by *something* even if we doubt that they are regulated by an attempt to track some objective ethical reality. Noncognitivists might be able to argue that ethical claims are regulated by higher-order conative attitudes, but this seems both ad hoc and a stretch. Let me now suggest, however, that my inferentialist version of expressivism has the resources for a similar but interestingly better answer to the relativist’s challenge.

Because, on my account, ethical claims express practical commitments, and practical commitments exhibit a different direction of fit with the world than theoretical commitments, they aren’t aimed at tracking some antecedently existing moral reality. This lack of regulation for the independent ethical facts explains why practical commitments fail to converge in the way that theoretical commitments tend to. And this explains the robust and recalcitrant diversity of ethical opinion.

Does this capture the right sort of disagreement to satisfy our intuitions? Notice that it is better than mere disagreement in attitude. At the level of expressions of first-personal *singular* practical commitments, we get at least a sort of robust *practical* conflict—as in Anscombe’s (1957) example of one who says, “I’m going to bed” and another who responds, “No, I’m going to stop you.”

However, one might still worry that this sort of practical conflict is too thin to capture all intuitively genuine moral disagreements. This is where it is helpful to bring in the idea that specifically ethical claims are the ones which express first person *plural* practical commitments of the form (M). With this, I think we can capture more genuine disagreement and in a more plausible way. The idea is to treat each person’s ethical claim as expressions of a first-personal plural commitment of the form “We all (rational agents) shall/shall not ϕ .” This provides an

even more robust practical sense in which there are genuine ethical disagreements. The intuition that ethical claims and disagreements are regulated by something, even if they are not regulated by an attempt to track an objective ethical reality, could, in this framework, then be captured by saying that they are regulated by the strictures of what counts as good higher-order practical reasons for a claim of the form “We all (rational agents) shall...”.

The fourth challenge was the *challenge of creeping minimalism and valid reasoning*, which came from reflection on quasi-realist accounts of the meaning of ethical claims. Traditional noncognitivists hold that ethical claims are (a) not truth-evaluable and (b) expressive of conative mental states rather than beliefs. Quasi-realist noncognitivism avoids the core problems of emotivism by dropping (a) and holding onto (b). However, this quasi-realist gambit generated a serious problem. Insofar as noncognitivism gains advantage over descriptivist accounts of the meaning of ethical claims because it completely adverts the need to posit ethical facts, it seems to require some general criterion for what distinguishes descriptive discourse from nondescriptive discourse—otherwise it risks succumbing to a sort of general quietism or idealism. But by “going minimalist” about notions traditionally thought to do the work of separating descriptive from nondescriptive discourse, quasi-realism incurs the obligation to come up with some Ersatz notion to do this work. Yet most of the same considerations that force quasi-realists to go minimalist about ‘truth’ also force minimalism about ostensibly ontologically committing psychological notions such as ‘belief’ or ‘representation’. This is the problem of creeping minimalism. And, as we saw, there seems to be little solace in retreating to technical psychological distinctions—as Gibbard and, following him, Dreier do—since it has not yet been shown that such distinctions both have ontological import and allow for reasonably construing ethical claims as not carrying ontological commitment. This generates a challenge for anyone who endorses a nondescriptivist account of the meaning of ethical claims: they need a plausible substantive account of what exactly it is about ethical discourse clearly that makes it nondescriptive that avoids the problem of creeping minimalism.

Obviously this challenge doesn't threaten descriptivist accounts since they can use any of the standard notions (e.g. 'truth', 'representation', 'belief') to make the relevant distinction and be happy with the fact that ethical discourse comes out on the descriptivist side of these standard ways of making the distinction. However, this dialectical situation is unstable since none of the descriptivist accounts seem to fully meet the first three challenges, and precisely because of the fact that they construe ethical discourse as descriptive of some independent ethical reality. So what one hopes for is either a descriptivist account that fully meets the first three challenges or a nondescriptivist account that can meet the fourth challenge while also meeting the first three challenges. We've just seen how my inferentialist version of expressivism can meet the first three challenges, but since it is a nondescriptivist account, this would be little progress if it couldn't also meet the fourth challenge. But I think it clearly can.

I'm inclined to think that real progress is made by deflating the ontological significance of terms such as "truth", "representation", and "belief". This is not based on a conviction of the incoherence of ontologically committing construals of these terms but instead on the observation that such construals tend to obscure rather than illuminate important metaethical issues. (If "true" means "correct description of reality" then we will need some other term "p-true" to mean "correct understanding of what to do". So why not just let "true" mean "correct" and then distinguish between modes of truth?)

But what makes ethical discourse nondescriptive, in my view, is precisely its practical character. Ethical claims do not ontologically commit one to the existence of an independent ethical fact; rather they are the expressions of a special sort of practical commitment, which lack ontological commitment because they exhibit a different direction of fit with the world from theoretical commitments. This is consistent with the minimalist construal of the traditional sorts of terms ('truth', 'representation', 'belief') that were thought to mark the descriptive-nondescriptive division; so there's no problem treating ethical claims as truth-evaluable

expressions of ethical representations or ethical beliefs, as long as these terms are minimalistically construed.⁷³ However, this view doesn't rest the distinction between descriptive and nondescriptive discourses on some technical psychological distinction whose ontological import is questionable—rather it rests it on the distinction between theoretical and practical reasoning.

One might worry that this move just rests the distinction between descriptive and nondescriptive discourses on another sort of technical distinction whose ontological import is equally questionable. But I think this isn't fair. Although the distinction between theoretical and practical reasoning is technically stated, the difference in reasoning that it marks is both intuitive and common. Of course, ordinary reasoners do not have explicit views about the direction of fit or the ontological import of their reasoning practices; however, by looking at the sorts of reasons that it seems natural or intuitive to give for different sorts of claims, two important points became clear. First, some claims—the ones that express doxastic commitments—demand theoretically sufficient evidence for their justification. By contrast, other claims—the ones that express practical commitments—demand practically sufficient reasons for their justification. Second, ethical claims seem to fall clearly on the latter side of this division. For, as we saw, in the ordinary practice surrounding the justification of an action, it seems sufficient to make an ethical claim as a reason for acting and it seems insufficient to give purely theoretical claims in support of an ethical claim. In the preferred idiom of the normative-role semantics sketched above, this Humean *cum* Moorean point is a reflection of the fact that ethical claims express practical rather than theoretical commitments.

⁷³ It's also consistent with other accounts of these notions that turn on the rejection of a realist account of truth but also shy away from full-fledged minimalism. In particular, in my (in preparation), I argue that an inferentialist expressivist can endorse the minimalist semantic claim that truth is not a robust sentence-world relation while denying the minimalist explanatory claim that the notion of truth is explanatorily impotent when it comes to explaining things such as semantic content, epistemic support, and logical connections. The key is to recognize (at least) two different kinds of truth (whose normative role can be articulated inferentially)—theoretical and practical truth. But the important point here is that my account of the expressive difference between ethical and matter-of-factual claims doesn't turn on a distinction made in terms of truth-evaluability.

This, I think, is what wins the inferentialist version of expressivism an advantage over its noncognitivist predecessors. For this is both a plausible and substantive account of what about ethical discourse makes it nondescriptive.

In this chapter, my preferred account of ethical claims has finally emerged in the wake of a careful consideration of the nature of the expression relation deployed in the core expressivist thesis—viz. that ethical claims express something with a different direction of fit with the world from descriptive claims. In the previous chapter, we saw that all contemporary noncognitivists cash out this idea in terms of what kind of mental state is conveyed by various claims (considered as utterances of a declarative sentence). In this chapter, we saw problems attaching to the two most prominent ways to understand this psychologistic model of expression_(p). Also, in this chapter, a new—normative—way to understand the core expressivist idea has emerged: ethical claims express a collective practical commitment of a particular form, while descriptive claims express theoretical commitments. This is the version of inferentialist expressivism that I want to recommend. In some ways, it piggy-backed on earlier expressivist's ability to meet the first two challenges. However, it adds significant resources to an expressivist answer to the third challenge. And we've just seen that the view is also well suited to meet the fourth challenge.

In the following chapters I want to switch the focus from the well-worn paths of metaethical debate to the weedy beginnings of metaepistemological debate, which has emerged due to the recent flourish of contextualist theories of knowledge attributions. In chapter 2, I will try to identify the motivation for and the problems with contextualism; in many important respects, this will mirror the discussion of descriptivist accounts of ethical claims in chapter 1, and three parallel challenges to an account of knowledge attributions will emerge. Then, in chapter 4, I

will continue the parallel by exploring the possibility of an expressivist account of knowledge attributions. I think most philosophers have simply assumed that only cognitivist accounts of knowledge attributions could be at all plausible, and because expressivism has traditionally been closely allied with noncognitivism, they have quickly dismissed if not completely ignored epistemic expressivism. But with my inferentialist version of ethical expressivism now on the table, there is room to separate the expressivist approach from noncognitivism, and room for a more plausible version of epistemic expressivism.

3 DESCRIPTIVIST ACCOUNTS OF EPISTEMIC CLAIMS

At the outset of this dissertation, I said that I have two principal aims: First, to develop a new account of the meaning of ethical claims that overcomes the drawbacks of previous attempted accounts. Second, to show how this metaethical account can be extended into an account of the meaning of epistemic claims in order to overcome drawbacks of previous attempted metaepistemological accounts. What I have to say in pursuit of the first aim is complete; it is now time to turn to the second aim. I want to pursue the second aim by way of a rough analogy. I think there is, at a high enough level of abstraction, an important analogy between the conceptual space of metaethical debate and the conceptual space of metaepistemological debate. And this analogy can be used to expose attractive regions of conceptual space that might otherwise be obscured.

In chapters 1, I considered descriptivist accounts of the meaning of ethical claims and found all of the prominent accounts lacking in at least some respect, which generated a series of challenges that I ultimately argued to be best met by my inferentialist expressivist account of ethical claims. In this chapter, I shall pursue the analogous task of considering descriptivist accounts of the meaning of epistemological claims (I shall focus on knowledge attributions).

What I find is that all of the prominent accounts⁷⁴ are lacking in at least some respect, which I again use to generate a series of challenges. In the following chapter, I shall then pursue a nondescriptivist alternative analogous to the expressivist view of ethical claims defended in chapter 2.

In considering descriptivist accounts of ethical claims in chapter 1, I divided the prominent accounts into three families: (i) nonnaturalism, (ii) naturalism, and (iii) relativism. This was a reflection of specific features of metaethics, and for the purpose of drawing out the major descriptivist accounts of epistemological claims, it proves helpful to pursue a different (although consistent) principle of division based on possible reactions to the phenomenon of apparent linguistic disagreement. In chapter 1, I outlined three possible descriptivist explanations of this phenomena. To use the terminology common in ethical debate these are: (a) universalism, (b) subject-relativism, and (c) speaker-relativism. Interestingly, recent metaepistemological debate has focused on cases of apparent linguistic disagreement in the use of knowledge claims, and the most prominent positions can be seen as correlatively epistemological species of each of these general approaches. They go under the names: (a') invariantism, (b') subject-sensitive invariantism, and (c') contextualism.⁷⁵ Naturally the problematic of metaepistemology is importantly different from that of metaethics, so in what follows I want to work through what I take to be the two most important aspects of that problematic in order to draw out the

⁷⁴ To be precise, the accounts I shall considered are offered as accounts of the truth-conditions of knowledge claims. I think the most natural way to understand these accounts and the way most participants to this debate understand these accounts is as accounts of the facts stated by knowledge claims. This is why I count them as forms of descriptivism. However, it is possible understand them in some other way—indeed, the sort of nondescriptivist account of knowledge claims that I articulate in the following chapter could be seen as one alternative interpretation of these truth-conditional accounts. Thanks to Ram Neta for pressing me to recognize this.

⁷⁵ This division of the metaepistemological landscape doesn't *exactly* line up with the division of the metaethical landscape (for example, it seems that there could be a noncontextualist but nonetheless speaker-relativist account of knowledge claims). This slight mismatch doesn't, however, affect the structure of the argument to follow, and I will ignore it in what follows. The reasoning for the names of the prominent metaepistemological positions will emerge more fully below.

difficulties I see with the three most prominent metaepistemological accounts. This will require a different organization of the discussion.

Specifically, in section I, I shall present three popular approaches to the semantics of knowledge attributions. In section II, I will consider what has been the de facto battleground for these three approaches—viz., a particular kind of skeptical paradox generated from the fact that our evidence for what we take ourselves to know generally doesn't rule out all counterpossibilities to this knowledge. In this battle, I think contextualism wins; however, the opponents of contextualism have rightfully stressed the fact that contextualism has counterintuitive results when it comes to the semantics of knowledge claims made in ordinary (i.e. nonskeptical) contexts. In section III, I shall highlight these problems with contextualism.

Thus, by the end of the chapter all three types of descriptivist accounts of the meaning of knowledge claims will appear problematic. And in the course of bringing this out I will attempt to crystallize three concrete challenges that any satisfactory account of the meaning of knowledge claims should be able to meet. These will then motivate and frame exploration of nondescriptivist alternatives in the following chapter.

I. THREE ACCOUNTS

Reflection on patterns of assertion and denial for knowledge claims in ordinary discourse reveals much apparent disagreement, which is not so plausibly explained in terms of a lack of information or a failure of rationality on one or the other parties' part. The examples of this are now familiar from the literature. When not much hangs on it, we'll claim to know that the bank is open on Saturday based on remembering being there a month ago on Saturday. In another context where a lot hangs on being right, we'll take the possibility that the bank has changed its

opening hours as a defeater to knowledge.⁷⁶ In normal contexts we will freely say that an acquaintance knows where his car is based on his remembering where he parked it, but in a context where it's very important to be right, the possibility that his car has in the meantime been stolen may lead us to deny the knowledge that we freely attributed in the normal context. In these cases, we have knowledge claims and their negations made in different sorts of contexts, and both seem correct.

Unger (1975, 1984) argues that the semantic situation here is analogous to that of attributions of 'flat'. It's common to call table surface flat although it has many microscopic bumps that might lead someone running a very sensitive experiment to deny its flatness. A golf course architect and a farmer might reach different conclusions about whether a particular piece of land is flat, and both can seem to be right. Unger argued that there are two competing semantic approaches one might take towards explaining this flexibility. First, one might be an *invariantist* and maintain that there is just one property of flatness that all claims of the form, "x is flat" are attributing to x—e.g. the property of having no bumps at all. Alternatively, Unger suggested, one might be a *contextualist* and contend that there is no unique property of flatness, but instead a family of properties—e.g., the property of being-level-enough-for-purpose-x, where 'x' is a variable—such that flatness-attributions have an implicit indexical, which gets filled in by contextual features.

Applying this distinction to knowledge attributions, Unger envisioned competing semantic accounts for knowledge attributions. The epistemic invariantist maintains that knowledge claims affirm and deny of a complex binary relation; schematically, it is the relation of '___ has an warranted true belief that___', where "warranted" is a technical term to mean whatever has to be added to true belief to make knowledge. The epistemic contextualist, on the other hand, maintains that knowledge claims affirm and deny a complex tertiary relation; schematically, it is

⁷⁶ This case is due to DeRose (1992).

the relation of ‘___has a true belief that___ which meets contextually determined epistemic standards___’.

However, as we can see from the metaethical distinction discussed in ch. 2 between speaker-relativism and subject-relativism, this contextualist position is importantly ambiguous. The ambiguity can be eliminated by specifying whose context is relevant to the meaning of a knowledge claim—the context of the subject *S* or the context of the person making the claim. This distinction collapses in first-person present tense knowledge claims; however, since they pull apart in other cases, it’s important to have separate terminology for them. Under the influence of its most prominent defenders—Cohen (1988, 1998), DeRose (1992, 1995), and Lewis (1996), the term “contextualism” has come to denote the view that knowledge claims are relative to the *speaker’s* context. And although the terminology is not as entrenched, the term *subject-subject-sensitive invariantism* has come to denote the view that knowledge claims are relative to the subject’s context.⁷⁷ The idea is that the relation attributed by knowledge attributions is invariant across context of attribution, but it is nonetheless a tertiary relation sensitive to features of the subject to whom knowledge is attributed.

Thus, one way to capture the three-way distinction between absolute invariantism, subject-sensitive invariantism, and contextualism is to notice that all three views could accept that

(1) *S* knows that *p*,

has the truth conditions given by,

(2) *S*’s true belief that *p* meets epistemic standards *e*,

but they will disagree on what determines the value of the variable *e*. The traditional invariantist holds that the value of *e* doesn’t vary. The subject-sensitive invariantist holds that the value of *e*

⁷⁷ This name is due to Stanley (2004, 2005). Hawthorne (2004) calls his very similar position “sensitive moderate invariantism”.

varies with features of S. And the contextualist holds that the value of e varies with features of the context in which (1) is affirmed or denied.⁷⁸

II. SKEPTICAL PARADOX

Of invariantism, subject-sensitive invariantism, and contextualism, which is the best account of the meaning of knowledge claims? One way this question is often approached is by investigating whether any of these accounts provides helpful resources for addressing skepticism.

There is more than one skeptical argument; however, the sort of argument that has played the most prominent role in the debate between invariantists, subject-sensitive invariantists, and contextualists is the argument from counterpossibilities inspired by Descartes. The basic idea is that for almost any proposition p that we take ourselves to know, we can generate a counterpossibility q , for which q is consistent with all of our evidence, it seems that we don't know that not- q , and, if we don't know that not- q then we also don't really know that p . More explicitly: where q is a counterpossibility consistent with all of my evidence and incompatible with p ,

⁷⁸ This is meant as a map of the three most prominent accounts of the meaning of knowledge claims, not as a complete map of the space of conceptual possibilities allowed by taking (1) to provide the truth conditions of (2). Most epistemologists including traditional invariantists would allow that the value of e varies with the value of p —that is the epistemic standards a belief must meet in order to count as knowledge depend on the content of the belief. However, we could imagine positions that deny this. In a different vein, MacFarlane (2005) and Richard (2004) also defend views according to which the truth-conditions of a knowledge attribution are relative to an epistemic-standards variable, but, for them, the variable is fixed neither invariantly, subject-sensitively, nor contextually. Rather they construe knowledge attributions as relative to the standards of the context from which the token attribution is *evaluated*. Thus, on their view, not only do knowledge attribution types not have fixed truth conditions, knowledge attribution tokens lack fixed truth conditions. As such, these views do not assimilate knowledge attributions to the standard semantics for implicit indexicals. Although I think such a view avoids some of the objections threatening invariantism, subject-sensitive invariantism, and contextualism, it does so at the high cost of abandoning the idea that token utterances of truth-evaluable claims have determinate truth conditions and endorsing relativism about truth. It's beyond the scope of this work to evaluate the merits of this move, but I doubt that it is forced upon us in our search for an adequate account of knowledge attributions.

- (P1) I don't know that not-q.
- (P2) If I don't know that not-q, then I don't know that p.
- (C) Thus, I don't know that p.⁷⁹

This argument, if sound, appears to undermine any knowledge not based on deductive entailment from known evidence.⁸⁰ It generates the following paradox: for (almost) any given p, which I would ordinarily claim to know without much hesitation, if the skeptic is clever enough, he can generate an argument where (P1) and (P2) both seem true, yet despite the apparent soundness of the argument it seems absurd for me to deny that my ordinary self-attributions of this knowledge are false. The reason (P1) will seem true is that the possibility that not-q is true is consistent with all of my evidence. The reason (P2) will seem true is that it is a consequence of a plausible closure principle, according to which if I know p, and p entails not-q, then I can come to know that not-q simply by deducing it from p. The reason (C) will seem false is that denying that I know p ostensibly contradicts all of the instances where I ordinarily and apparently correctly attribute such knowledge to myself. That is, the following claims seem individually true but mutually inconsistent:

- (i) For almost any propositions p, there is a skeptical counter-possibility q consistent with everything I know, and I do not know q not to obtain.

⁷⁹ This argument is discussed in detail in DeRose (1995).

⁸⁰ Does this argument really best capture the skeptical worries that troubled Descartes or others who think that skeptics have an argument that needs to be answered? One might reasonably worry that no skeptic worth answering would *assume* that (P1) is true. Perhaps skeptical considerations can force us to doubt whether we know that we're not systematically deceived by an evil genius, but should we allow the skeptic to *assume* that we don't know that we aren't so deceived? Perhaps not. Klein (1981) and Pryor (2000) propose alternative reconstructions of the skeptical reasoning that don't rest on premises like (P1) (and they attempt to refute these more subtle skeptical arguments). Here, however, I'll stick with the argument as formulated in the text because it lets me best reconstruct the reasons that have played a central role in the debate between contextualists, invariantists, and subject-sensitive invariantists. See Neta (2004) for a discussion of how a contextualist can handle the sort of skeptical arguments developed by Klein and Pryor.

- (ii) For any p and any q, if I know that p, and p entails not-q, then I can come to know not-q by deducing it from p.
- (iii) For many propositions p, I know that p.

Call this the *AFC1-paradox*.⁸¹

As with any ostensible paradox, there are as many “hard-line” responses as there are independently plausible but mutually inconsistent claims comprising the paradox. In this case, one could deny (i), (ii), or (iii). However, each of these hard-line responses will be only as convincing as the story its proponents can come up with for why the claim they deny seems true.

The hard-line response that received much attention in the 70’s and 80’s was the possibility of denying the closure principle embodied in (ii).⁸² I don’t have space to get into this debate in any detail, but I think the upshot is that the closure principle can use some sharpening to avoid counterexamples. For example, it might help to stipulate that I maintain my knowledge of p throughout my deduction of not-q and that this deduction is competent in that it doesn’t proceed through some false lemma. However, the core idea is extremely plausible and shouldn’t be denied. This leaves (i) and (iii).

Unger’s original invariantism was a skeptical position, which means that he would deny (iii).⁸³ He attempted to explain why we mistakenly take (iii) to be true by arguing that although most knowledge attributions are literally false, they are warrantably assertable because they are close enough to being true to convey the relevant information, which is that the subject stands in

⁸¹ For “paradox from 1st personal argument from skeptical counter possibilities”. Below I shall introduce and discuss a third-personal analog to this paradox.

⁸² See, for instance, Dretske (1970), Stine (1976), Nozick (1981), and Feldman (1985).

⁸³ Unger(1975:50-54). See also Stroud(1984: 55-82) for a sympathetic presentation of the skeptical position and a version of the argument to follow.

a very good epistemic position vis-à-vis p. By way of comparison, Unger argues that the ordinary term ‘flat’ is an “absolute” term, meaning that when a surface is flat it is impossible for another surface to be flatter. Then he points out that almost no surface is really flat, but many surfaces are flat enough for whatever one’s purposes are when one says that they are flat. This means that, in calling a surface flat, one is almost always asserting a false proposition. But Unger argued that there is no problem in doing this, indeed he suggested that doing so is a more convenient way for ordinary speakers in ordinary circumstances to convey the information that the relevant surface is flat enough for the purposes at hand than asserting the more complex true proposition. This is one way of construing ordinary flatness attributions as warrantably asserted even though false.

So, a skeptical invariantist like Unger maintains that (iii) is false, because, on the skeptical view, it’s usually false to attribute knowledge. However, he argues that (iii) is close enough to being true that many knowledge attributions are warrantably assertable. Below I shall consider nonskeptical versions of invariantism that start from a denial of (i), but first I want to argue that skeptical invariantism has two serious drawbacks.⁸⁴

First, the position requires an implausible error theory about all of the ordinary attributions of knowledge that we make all of the time. In this sense, it is like Mackie’s ethical skepticism, which in chapter 1 I argued faces a huge (and arguably insurmountable) explanatory burden of accounting for why we (putatively) mistakenly make all of the ordinary ethical value claims that we make. In both cases, one is reminded of Edwards’ (1955) example of the wacko philosophical challenge to traditional views about what “physician” means.⁸⁵ A crazed philosopher of medicine might propose that, contrary to orthodoxy, one of the necessary conditions for being a physician is that one be able to cure any conceivable illness. This theory seems patently false, precisely because we ordinarily call many people physicians who cannot cure every conceivable illness. It

⁸⁴ The following criticism of skeptical invariantism is indebted to DeRose(1995/1999: 210-213).

⁸⁵ This is cited by Stroud (1984: 40).

wouldn't help sympathizers of the theory to argue that these claims only seem true because while literally false they are warrantably assertable. For, as DeRose puts the point,

It's no doubt largely in virtue of...facts [about ordinary usage]...that the traditional view, rather than the conjecture of our crazed philosopher, is true of our language. (The correctness of the traditional view largely *consists* in such facts...that the peculiar theory implicates us in systematic and widespread falsehood in our speech and thought involving 'physician' is a (constitutive and evidential) strike against the theory that proves quite decisive.(1995/1999: 213)

The situation seems to be similar with knowledge claims: a theory that posits widespread error in the ordinary usage of a term requires what never seems to be forthcoming—viz. a highly compelling error-theory. For, as I urged in connection with ethical terms, what our words mean, at least in an important sense, is partially constituted by the way we as a linguistic community use them.⁸⁶ This doesn't mean that there cannot be terms that are widely misused; however, any account of the meaning of some term 'x', which posits widespread error in the ordinary usage of x, will need a very good argument for the claim that the account is an account of the *ordinary* term 'x' and not some esoteric philosophical cousin 'x'. Insofar as this is lacking, the account fails to meet what I want to call *Austin's challenge*.⁸⁷ And it seems that skeptical invariantism lacks such an argument and thus fails to meet Austin's challenge in the extreme.

The second drawback to skeptical invariantism is that appealing to warranted assertability to explain why we think most of our ordinary knowledge claims are true when they are literally false undermines the skeptical invariantist's resources for explaining the alleged truth of (i). If we are indeed so ignorant of the literal content of our knowledge claims that we

⁸⁶ I think it is usually a problem with an account of the meaning of a term if it requires us to attribute and then explain widespread error in the ordinary use of the term; however, there are cases where I suspect the explanatory burden can be met. For example, the term 'witch' was freely attributed to many people during the Salem witch trials in a way that was erroneous. However, if one considers all uses of the term 'witch' in modern English, this posit of linguistic error seems mild in comparison to the skeptical invariantist's hypothesis. Moreover, I suspect that the practical consequences of calling someone a witch in 18th century New England can go a long way to explaining the error.

⁸⁷ Named after Austin's (1962) suggestion that philosophical skepticism results from a verbal fallacy of, in effect, construing the meaning of 'knows' in such a way that it completely disconnects from ordinary usage of this term.

think that many knowledge attributions are true when they are really false, the intuition that instances of (P1) is true surely cannot count as very good evidence for holding onto (i). Why privilege these intuitions while eschewing so many similar intuitions? It looks ad hoc for the skeptical invariantist to argue that, just in the case of instances of (P1), our intuitions track the literal content and truth of our knowledge claims. I want to call the challenge of providing an at least coherent resolution of the paradox the *challenge of the paradox*. And I think the fact that the skeptical invariantists reasoning for rejecting (iii) undermines his reasoning for accepting (i) indicates that the position doesn't really have a stable answer to the challenge of the paradox.

I don't think these drawbacks *refute* the skeptical invariantist attempt to answer the AFC1-paradox by rejecting (iii), but they certainly motivate exploration of other options in hopes of something more satisfactory.

There are also non-skeptical invariantists that pursues a converse strategy in response to the AFC1-paradox. They reject (i) arguing that for any skeptical counter-possibility q, it's actually literally true to say, "I know that not-q," but it seems false because it is not usually warrantably assertable because, e.g., it is highly irrelevant or misleading. For example, Rysiew (2001) reminds us of Grice's (1989) distinction between "sentence-meaning", i.e. the literal propositional content of the sentences uttered, and "speaker-meaning", i.e. the propositions the speaker intends to convey by uttering the sentence; then he argues that while the literal content of (P1) is true, almost any utterance of it would convey a false proposition. This is because he thinks the literal content of a knowledge attribution,

(1) S knows that p,

is given by the truth conditions,

(2') S has a true belief that p for which he can rule out the *relevant* alternatives,

where “relevant alternatives” are the counter-possibilities to *p* that a normal human would think were likely; however, he thinks that the conversational pragmatics surrounding knowledge attributions are such that the speaker-meaning of (1) is typically,

(3') S has a true belief that *p* for which he can rule out all *salient* alternatives.

where the “salient alternatives” are potentially different from “relevant alternatives” in that they are counterpossibilities to *p* that happen to be salient to the participants to the conversation in which (1) is uttered (whether or not these counterpossibilities would be thought to be likely by normal humans). Thus, since the skeptical counterpossibilities are designed precisely so that we cannot rule them out, and an utterance of (P1) makes such counterpossibilities salient, it would always be misleading to deny (P1). This is because the speaker-meaning of such an utterance would be false. However, Rysiew thinks that to deny (P1) would nonetheless not be to utter a literal falsehood. This is because the sentence-meaning would be true. While more plausible than the skeptical invariantist’s response to the AFC1-paradox, I think this view has two drawbacks.

First, it also involves an implausible error-theory, although of an importantly different sort from the skeptical invariantist. The moderate invariantist doesn’t have to attribute linguistic error to ordinary speakers in their widespread usage of knowledge claims; however, he has to attribute linguistic error to people who have been puzzled by the AFC1-paradox. For they thought that (i) is true because instances of (P1) seem true in the context in which they are presented for the skeptical argument, but according to Rysiew, this is because they are mistaking the typical speaker-meaning of “S knows that *q*”, i.e. (3') for its sentence-meaning of these sentences, i.e. (2'). However, I think we can grant Rysiew that his assumption “that speakers often confound the proposition expressed with the proposition imparted”(Ibid.: 483) is true of ordinary speakers who aren't familiar with the distinction between sentence-meaning and speaker-meaning or the nature of skeptical arguments. Nonetheless I think we may still

reasonably doubt that the philosophers who have discussed the skeptical arguments and paradoxes such as the AFC1-paradox can be plausibly attributed the error that generates the paradox.⁸⁸ To be sure, this involves only a relatively small posit of linguistic error, and so the moderate invariantist will have a much more satisfactory response to Austin's challenge than the skeptical invariantist. But we might hope for something better nonetheless.

The second drawback of the view is more serious. In a sense, Rysiew's explanation doesn't go far enough. For even if it were plausible to think that those people (philosophers) who are perplexed by the AFC1-paradox are mistaking (2') for (3'), this only explains why they are unwilling to *assert* that S knows that q. It doesn't explain why they are positively willing to *deny* that S knows that q. But the paradox turns on our thinking (mistakenly in Rysiew's view) that (P1) is *true*. As MacFarlane puts the point,

...although worries about misleading implicatures may be good reasons to refrain from asserting something, they aren't good reasons to assert its negation. Before Cal has played any games, I will refrain from asserting (truly) that Cal has won all of its games so far this season, because my doing so would misleadingly imply that Cal has played at least one game already. But these considerations do not give me any reason to assert that Cal has *not* won all of its games so far this season. Similarly, even if Rysiew's story can explain why it would be rational for me to refrain from saying that I know, it cannot explain why I should say that I *don't* know. (2005: p. 14 of typescript)

But, since the skeptical paradox turns on taking (P1) to be true rather than merely being reluctant to deny (P1), Rysiew's explanation won't work. As far as I can tell, Rysiew's account is the best available non-skeptical invariantist attempt to answer the AFC1-paradox, but these two problems indicate that such an account doesn't fully overcome the challenge of the paradox.

I don't mean to claim that the problems we have seen for the skeptical and non-skeptical invariantist responses to the AFC1-paradox provide knock-down arguments against

⁸⁸ To be fair, I should mention that Rysiew concedes: "while I think the foregoing account of P1's plausibility is correct as far as it goes, I doubt that it's the whole story. For while this account might apply to naïve speakers confronting the hypothesis that they are BIVs for the first time, it has the consequence that no one who is both cognizant of the relevant distinctions...and adept at applying them is going to believe that P1 is true. But I don't want to claim this" (Ibid.: 503). However, his subsequent diagnosis of why those cognizant of the relevant distinctions and adept at applying them might think that P1 is true is difficult to follow and unconvincing.

invariantism. However, it seems that both positions fail to provide a stable response to the paradox, and skeptical invariantism fails to connect appropriately with ordinary epistemological discourse. These are problems that we should hope to avoid. I think they arise from the “hard-line” nature of the responses to the paradox, so in the following sections I want to consider the prospects of a “soft-line” response, which attempts to finesse the paradox.

Contextualists and subject-subject-sensitive invariantists attempt to finesse the AFC1-paradox by accepting the soundness of the skeptical argument in the context and practical situation in which it is presented, but then arguing that the conclusion (C) is compatible with a commitment to the legitimacy of many ordinary self-attributions of knowledge that p made in different contexts and practical situations.

Specifically, the contextualist will argue that the standards a true belief must meet to be truly called knowledge depend on the context in which the knowledge attribution is made, and, if this is true, it is perfectly consistent to think that (C) is true in the context where the skeptical argument is presented, while maintaining that apparently conflicting affirmations of the knowledge in everyday contexts don’t really conflict because they are made in a different context and thus indexed to different standards. And the subject-subject-sensitive invariantist will argue that since the standards a true belief must meet in order for it to be knowledge depend on the practical situation of the subject, and, if this is true, it is perfectly consistent to think that (C) is true in the practical situation of the subject thinking through the skeptical argument, while maintaining that apparently conflicting affirmations of knowledge in everyday contexts don’t really conflict because the nature of the subject’s practical situation and correlatively what epistemic standards one of his true beliefs must meet in order to be knowledge are different.

The contextualist and subject-sensitive invariantist responses to the AFC1-paradox are very similar. Are they equally plausible? Arguably the subject-sensitive invariantist response just

outlined piggy-backs on the contextualist response in an artificially convincing way. To see this, notice that the fact that the context of use of the relevant knowledge attributions and the subject's practical situation are basically identical is an artifact of the fact that the skeptical argument is presented in the first person. However, although there has been some debate in the epistemic literature about whether Descartes' skeptical reasoning is essentially first-personal, there does not seem to be any reason why a parallel skeptical argument in the third person should be any less troubling than the first-personal skeptical argument presented above. That is, we can generate a similar third-personal paradox out of the following skeptical argument:

- (P1') S doesn't know that not-q,
- (P2') If S doesn't know that not-q, then S doesn't know that p,
- (C') Thus, S doesn't know that p,

where S is someone in a context and practical situation removed from the context and practical situation of the presentation of the argument. This change doesn't seem to make it any less plausible to think

- (i') For almost any propositions p, there are skeptical counter-possibilities q consistent with everything S knows, that S does not know not to obtain.
- (ii') For any p and any q, if S knows that p and p entails not-q, then S can come to know not-q by deducing it from p.
- (iii') For many propositions p, S know that p.

Again, however, these initially seem to be individually plausible but mutually inconsistent—call this the *AFC3-paradox*. What's important to notice about the AFC3-paradox is that while the contextualist response to the AFC1-paradox applies *mutatis mutandis* to this paradox, the subject-sensitive invariantist response does not. This is because the paradox arises even when the practical situation of S is held constant between the context in which the skeptical argument

is presented and (C') seems true and more ordinary contexts where it also seems true to affirm that S knows that p, so subject-sensitive invariantism fails to explain away the ostensible contradiction.

The way that subject-subject-sensitive invariantists typically respond to this objection is by arguing that there is an important sense in which we are prone to error in our use of third person knowledge claims. For instance, Hawthorne (2004) argues that in making third personal knowledge claims we tend to project our own practical situation onto the subject. He thinks this is true for a combination of reasons. For one thing, we won't always be aware of the subject's practical situation, so the best we can do is, so to speak, fill in the gaps with our own practical situation. But more importantly, Hawthorne argues that we often mistakenly judge the probability of various counterpossibilities to the subject's knowledge, thus making her belief seem more precarious than it really is.

So subject-sensitive invariantism requires the posit of a certain sort of systematic semantic blindness in the use of third-person knowledge claims. With this posit, subject-sensitive invariantism opens itself up to a series of problems.

First, insofar as the putative error is a systematic and widespread error in the ordinary use of knowledge claims, a version of Austin's challenge threatens subject-sensitive invariantism. Whenever a semantic theory posits widespread and systematic error in the ordinary use of a term, we should worry that the theory has illicitly disengaged from the ordinary term. Of course, this worry is not always apt, but it is incumbent on proponents of the theory to show that is not apt by providing a plausible explanation of the error. Hawthorne attempts an explanation by appeal to what psychologists call the "availability heuristic".⁸⁹ The idea is that people typically judge the probability of some proposition on the basis of how easy it is to imagine its being true. Since it's typically easier to imagine propositions being true that are similar to propositions that

⁸⁹ Hawthorne cites Slovic *et al.* (1980).

have often been true in the past, this heuristic for determining probability usually works in a roughly accurate way. However, sometimes factors irrelevant to the probability of a proposition can affect how easy it is to imagine its being true. For instance, watching a movie with a plane crash can make this possibility especially vivid and thus easier to imagine. This is why Hawthorne claims that “when a certain scenario is made vivid, the perceived risk of that scenario may rise dramatically”(2004: 164). Applying this to the epistemic case, notice that watching the *Matrix* can make the possibility of being a brain in a vat vivid and easier to imagine, although not any more probable. According to the availability heuristic, we will tend to overestimate the probability of this scenario when the movie is fresh to mind. And Hawthorne thinks this explains why we mistakenly and systematically deny knowledge to people in other contexts when skeptical scenarios are fresh to our minds. However, I think this explanation opens the subject-sensitive invariantist up to two further problems.

Second, when it comes to the AFC3-paradox, the subject-sensitive invariantist’s response will, in effect, have to be analogous to the non-skeptical invariantist’s response to the AFC1-paradox. One problem we saw above with that response is that it requires the implausible posit of a systematic sort of error not to ordinary speakers but to philosophers who have been troubled by skeptical paradoxes. Likewise, despite Hawthorne’s attempted explanation of mistaken intuitions about third-person knowledge claims, I think subject-sensitive invariantism faces this problem. For notice that even if Hawthorne is right that ordinary speakers *tend* to overestimate the probability of skeptical scenarios when they are made salient and thus more imaginable, philosophers whose job it is to think clearheadedly about things like the probability of skeptical scenarios should be able to buck this tendency. That is, even though skeptical scenarios become more salient and thus more imaginable in the epistemology seminar, most philosophers would not therefore overestimate their probability. But this doesn’t change the fact that when in the grip of skeptical discussions there is some pull to deny that people in other contexts have knowledge. This is why the AFC3-paradox seems like a paradox even when we are

thinking clearheadedly about the probability of skeptical scenarios. Thus, although Hawthorne's explanation seems to work for epistemological discourse after casual viewers watch the *Matrix*, it fails to work for philosophers who are puzzled by the AFC3-paradox. This means that subject-sensitive invariantism does not completely meet the challenge of the paradox.

Third, Hawthorne's use of the data in favor of the availability heuristic is mistaken. As Cohen (2005) has pointed out, although use of the availability heuristic can lead to overestimation of the probability of certain possibilities, the original data indicated that it can also lead to the underestimation of the probability of other possibilities. When asked, subjects tended to *underestimate* the probability of possibilities that are perfectly easily to imagine (and salient due to the question) but *unspectacular*. For instance, while people will overestimate the possibility of plane crashes and homicide when these possibilities are salient, they will underestimate the possibility of being in a car accident or getting cancer when these possibilities are salient. This indicates that it is not the ease with which we imagine possibility that tends to affect our judgment of its probability, but rather its spectacularity or vividness. Applying this to the epistemological case, this means that Hawthorne's explanation makes sense of why, after considering the possibility that one is in the *Matrix*, one tends to deny knowledge to people in other contexts, but it doesn't make sense of why, after considering the possibility that a friend's car has been stolen, one tends to deny that the friend knows where his car is parked. But both spectacular and unspectacular skeptical possibilities can seem to undermine knowledge claims. I think this problem reveals Hawthorne's attempted explanation of the semantic blindness posited by subject-sensitive invariantism to be less than fully convincing, and so the burden remains to explain the error. Until it is satisfied, the subject-sensitive invariantist's account will fail to meet what I shall call the *challenge of unexplained semantic blindness*.

These problems are not knock-down arguments against subject-sensitive invariantism. However, I think they do motivate consideration of alternatives. Specifically, I think

contextualists are right to suggest that their unified approach to the AFC-paradoxes is more attractive than the mixed strategy required by subject-sensitive invariantism. Nonetheless, I think contextualism has its own problems resulting both from its response to the paradoxes and from its treatment of ordinary epistemic discourse. In the next section, I shall bring these to light.

III. ORDINARY DISCOURSE

It may seem that contextualism is the clear benefactor of our reflections on skeptical paradox. For contextualism clearly provides a way to overcome the versions of Austin's challenge facing skeptical and non-skeptical forms of invariantism. The contextualist explains our relative reluctance to attribute knowledge in skeptical contexts to the high epistemic standards of these contexts and our relative willingness to attribute knowledge in other contexts to the low standards of these contexts; and because these standards are construed as part of the truth-conditions of the relevant knowledge attributions there is no need to appeal to notions such as Unger's warrantably assertible falsehoods or Rysiew's false speaker meanings attached to true sentence meanings. Moreover contextualism provides a way to answer the AFC₁-paradoxes with the following advantages. First, unlike the non-skeptical invariantist's response to the paradox, the contextualist response doesn't turn on attributing linguistic error to philosophers who feel the pull of the skeptical arguments. Second, the contextualist response to the AFC₁-paradox can be extended in a natural way to the AFC₃-paradox, thus avoiding the implausible error-theory required by sensitive-invariantism. It's for these reasons that, I think contextualism is the best of the descriptivist responses to these skeptical paradoxes.

Nonetheless, I also think that contextualism faces its own (at least *prima facie*) problem due to its proposed resolution of the AFC-paradoxes. *Prima facie*, it looks as if contextualism proves too much in that it provides a general recipe for dispelling *any* paradox. The recipe is this: posit

hidden deixis in the truth conditions of the claims that, paradoxically, seem both individually true and collectively inconsistent, and argue that one of the claims is merely true when tokened in the specific context in which the paradox is stated but not true when tokened in ordinary contexts because the value of the deictic element varies in precisely the right way across the contexts. Because of this, the contextualist answer to the AFC-paradoxes is likely to appear ad hoc.

Now, how threatening this problem is depends on how well contextualists can independently motivate their posit of hidden deixis in knowledge attributions. And initially, it seems that they have some plausible evidence in their favor. For, as I mentioned above, there are many examples of ordinary knowledge claims that we think are true when uttered in one context but false when uttered in another context despite the fact that none of the traditional components of knowledge—justification, belief, and truth—have changed. Nevertheless, difficulties for the contextualist’s posit of hidden deixis aren’t hard to find. For instance, it has become standard to object that even if contextualism captures the linguistic data around base-level knowledge claims, it does so at the cost of positively disrespecting the linguistic data around metalinguistic evaluations of knowledge claims. DeRose invited this objection when he argued it to be an advantage of contextualism that it does *not* have to countenance the following absurd dialogue:

A: Is that a zebra?

B: Yes, it is a zebra.

A: But can you rule out the possibility that it is a cleverly painted mule?

B: No, I can’t.

A: So, you admit you didn’t know that it was a zebra?

B: No, I did know then that it was a zebra. But after your question, I no longer know.⁹⁰

The problem is that his reasoning behind the contention that contextualism doesn't have to treat the dialogue as legitimate is highly suspicious; he writes:

If in the context of the conversation the possibility of painted mules has been mentioned, and *if* the mere mention of this possibility has an effect on the conditions under which someone can be truly said to "know," then any use of "know" (or its past tense) is so affected, even a use in which one describes one's past condition. B cannot truly say, "I did know then that it was a zebra" ...[but] *can* say, "My previous knowledge claim was true," ...But saying these things would have a point only if one were interested in the truth-value of the earlier claim, rather than in the question of whether in the *present* contextually determined sense one knew and knows, or didn't and doesn't.(1992: 925)

So, according to this, the contextualist need not accept B's final sentence as anything less than absurd because the meaning of 'knows' in it is relative to the new epistemic standards implemented after the mentioning of the counterpossibility—assuming, of course, that this mentioning the counterpossibility does in fact raise the standards. Nonetheless, DeRose maintains that B also could have correctly said: "My previous knowledge claim was true."

Let's imagine that B actually did previously claim to know that the animal is a zebra. Can DeRose's account of what's going here be right? Several critics have suggested that it cannot be. For instance, Rosenberg writes,

DeRose ... holds that B *can* say—presumably, *correctly* say—"My previous knowledge-claim was true"... But I do not think that this is correct, For what *was* B's previous knowledge-claim?...The most natural answer, surely, is that he asserted *that he (then and there) knew that the animal in view was a zebra*; but if that is what B asserted, then what he asserted was *false*. As we have just seen, both DeRose and I agree that ... B did *not* know at the beginning of his conversation with A that the animal in view was a zebra. And if this is so then B cannot now *correctly* say that his previous knowledge claim was true. It wasn't.(2002: 164)⁹¹

Is this a reasonable objection? It certainly seems strange, assuming that the content of,

⁹⁰ This dialogue is originally due to Yougrau (1983) and was meant as an objection to Dretske's (1970) relevant alternatives account of knowledge claims.

⁹¹ The same point has been pressed by Hawthorne (2004) and MacFarlane (2005).

(4) My previous knowledge claim,
in the later context, is that,

(5) B (then and there) knew that the animal in view was a zebra,
that B could, in the later context, correctly say something with the content,

(6) (4) is true,
while it would be false, in that context, for him to say something with the content of,

(7) B (then and there) knew that the animal in view was a zebra.

If this a consequence of contextualism, then it is a highly counterintuitive consequence.

Now, it's worth noticing on the contextualist's behalf that there is a way out of this problem (although I think it points the way to a deeper problem). The assumption that the content of (4) is given by (5) is, by the contextualist's lights, ambiguous, since (5) could mean either,

(5') B's true belief that the animal in view was a zebra meets epistemic standards e' ,
or

(5'') B's true belief that the animal in view was a zebra meets epistemic standards e'' ,
where e' are the epistemic standards in the original context and e'' are the epistemic standards in the new context instated by the mention of a counterpossibility. Rosenberg's objection only goes through if (4) means (5'')—for even in the later context, a claim with the content of (5') would be true. Therefore, if in the later context (4) means (5') and (6) means (5''), then it is not inconsistent for the contextualist to count (6) as true but (7) as false. Rosenberg and the contextualist agree that (7) means (5''), so the point of disagreement is whether (4) means (5') or (5'').

However, with everything set out so explicitly at this point in the discussion, a sympathizer for contextualism might be tempted to simply concede that DeRose was mistaken to think that

(4) as uttered in the later context has truth conditions given by (5''). For although Rosenberg is probably right about the content of (4) as it is uttered in the later context, it is not clear that that concession undermines the contextualist's contention that B's original (imagined) knowledge claim *uttered earlier in the dialogue* has truth conditions given by (5'); and as long as the earlier claim has the truth conditions given by (5') and the later claim has the truth conditions given by (5''), contextualism is vindicated.

The situation here is tricky because the contextualist has to hold that in the later context, it becomes impossible for one to say that (5') is true by either uttering the object-level knowledge claim or by ascribing truth to the previously uttered object-level claim. It is as if one wrote a note that says, "It's hot here," while in a hot room and then took it outside where it is cold. It would be false to write an identical note, and it would be false to point to the previously written note and say "What that says is true." However, contextualism doesn't have the implication that it's impossible to say what one should be able to say. Just as one could write a new note that says, "It *was* hot in *there*" to express a true proposition, the contextualist can make his epistemic claim directly in terms of the truth conditions by saying, "While (5'') is false, (5') is true." In other words, the key for the contextualist is to mark out the distinction between knowledge attributions made in one context and evaluated in another context not in terms of a difference in the content between "S knows that p" and "It's true that S knows that p" but, rather, in terms of the truth conditions of the knowledge attributions made in different contexts.

Although I think the upshot here is that Rosenberg's objection fails to directly undermine contextualism, the discussion points the way to a deeper problem for contextualism about the dialectical connections between our knowledge claims. By treating the content of knowledge attributions as context sensitive, the contextualist loses the ability to account in a natural way for the dialectical connections between knowledge attributions. I suspect the reason (5'') seems to be a better candidate for spelling out the content of (4) than (5') is that we typically say things

like “My previous knowledge claim was true” or “My previous knowledge claim was false,” with the intention of, respectively, *maintaining* or *changing* our commitment to the claim. However, this is true of object-level knowledge claims too; we often reaffirm or deny the knowledge previously claimed with the intention of maintaining or changing our commitment to the claim. Indeed, this seems to be something that the contextualist should explicitly accept given that he attempts to answer the AFC-*paradoxes*. For surely part of the reason why the paradoxes seem to be paradoxes is that at the end of some skeptical arguments we’re inclined to deny knowledge that, before the argument, we asserted—and this seems to be a *change of view*. But it couldn’t be a change in view unless the two knowledge claims are dialectically connected. (I shall come back to this issue below.)

Bracketing skeptical paradoxes for a moment, we should also note that there are perfectly ordinary cases where two tokens of a knowledge attribution uttered in different contexts seem to stand in agreement and, correlatively, cases where a token affirmation and its apparent denial uttered in different contexts seem to stand in conflict despite the fact that, according to contextualism, there are different epistemic standards operative in the two contexts. For example, consider a case where Sharon and Chris are visiting their son Charlie in a different city. The morning of the flight, Chris asks Sharon, “Do you know whether Charlie will pick us up from the airport?” Sharon replies, “Yes, I talked to him yesterday and gave him our flight details.” We can assume that nothing serious hangs on whether Charlie will pick them up, and if he doesn’t they could take the train into town. But having talked to him Sharon takes herself to know that he will pick them up; and, assuming that it’s true that Charlie will pick them up, Sharon’s (implicit) knowledge claim seems true. Let’s imagine, however, that later Chris takes a call from a business associate proposing an important conference call with a client at a time that Chris could make only if Charlie does indeed pick them up from the airport (because the train takes too long). Chris tells Sharon of the proposal and continues, “Are you sure Charlie will pick us up? Perhaps he forgot the flight time.” On the contextualist view, it seems like the standards for

knowledge have been raised. Nonetheless, there seem to be at least two ways this case could go forward. Sharon might say, “I thought I knew, but I guess it is possible that he will forget the flight time, so I don’t know for sure. Maybe you should give him a quick call to confirm before you schedule the conference call.” Or she might say, “He promised to be there at 5:30, so he won’t forget the flight time, so I know he’ll be there. Go ahead and schedule the conference call.”⁹²

From the contextualist’s point of view, the standards for knowledge seem to be different between the context of Sharon’s first (implicit) knowledge claim and the context of her (possible) later knowledge claim—for the question of whether Charlie will pick them up has become more important due to the proposed conference call. However, in both continuations of the case, the later knowledge claim seems to manifest a dialectical connection to the earlier knowledge claim. Sharon either wants to deny knowledge and thereby *change* her earlier view that Charlie has knowledge or she wants to attribute knowledge and thereby *reaffirm* her earlier view. But if each of these knowledge claims are semantically relative to the standards operative in their context of utterance and these standards are different in the earlier and later contexts, then contextualism cannot capture these intuitive dialectical connections. This is because the context sensitivity imputed by the contextualist undermines the way in which two tokens of the same knowledge attribution type or a knowledge attribution token and its correlated disattribution, as uttered in different contexts with different epistemic standards, could express, respectively, agreement or opposition. However, as the case indicates, it often seems as if we do use knowledge claims in just this way. This is, in effect, a second-order version of the challenge of unexplained semantic blindness. For although the contextualist does not have to posit unexplained semantic blindness in the ordinary use of knowledge claims, he does have to posit it

⁹² This example and the objection I glean from it are taken from my (forthcoming).

at the level of inter-context comparisons of the dialectical connections between knowledge claims.

Does this problem generate a fatal objection to contextualism? Problematic intuitions never generate fatal objections to a theory because defenders of theory can always deny the accuracy of the intuitions. Of course, such a move incurs a cost: plausibility will be saved only if the defenders of the theory can strike a nice balance between downplaying the robustness of the intuitions and coming up with a plausible error theory for why we have mistaken intuitions. In DeRose (2005), we see just such an attempt to downplay the robustness of the sort of problematic intuitions we've just considered. He considers a case with two contexts, LOW and HIGH, where knowledge of a particular proposition is attributed to a particular subject in LOW but denied in HIGH, and writes,

When we consider the affirmation in LOW together with the denial in HIGH, is it clear that these two claims contradict one another? So far from that being obvious to me, I've always been fairly strongly inclined to think they do *not* contradict one another. But that's just me, and it is perhaps no accident that I became a contextualist. But when I have presented such cases to large enough groups of students, and then asked about whether the claims contradict one another, I have always found both positive and negative answers strongly represented."(2005: 194)

His suggestion is that in light of the varying intuitions about the dialectical status of the claims we shouldn't take them too seriously or think it a major problem with contextualism if it disrespects them.

I doubt that the intuitions behind the dialectical intuitions problem could be as weak as DeRose suggests. For notice that the contextualist's response to the AFC-paradoxes requires precisely the revelation that the conclusion of skeptical arguments is *despite years of it seeming otherwise to epistemologists* consistent with attributions to S of knowledge that *o* made in other nonskeptical contexts. However, if the intuition of dialectical conflict between cross-context

knowledge claims were really as weak as DeRose suggests, no one would be very exercised about the AFC-paradoxes.⁹³

Cohen (1999: 77-78) pursues a partners-in-crime response to this sort of objection. He argues that the same sort of semantic blindness among ordinary speakers posited by the contextualist attaches to ordinary semantic intuitions about ‘flat’, and yet most philosophers agree that ‘flat’ is a context sensitive term. The idea is that we could generate a case where, in a low-standards context someone calls some surface flat, but then the standards are raised and the person is asked whether she still thinks the surface is flat. Presumably, she might either change or reaffirm her view. Yet that doesn’t undermine contextualism about ‘flat’. Does this analogy undermine the objection? Well, it depends on how similar ‘knows’ and ‘flat’ are in their semantics.

Stanley (2004) argues that there is a significant disanalogy between the context sensitivity of terms such as ‘tall’ or ‘flat’ and the context sensitivity posited by the contextualist for ‘knows’. Specifically, ‘tall’ and ‘flat’, unlike ‘knows’, seem to be “gradable”, meaning that they come in degrees, as indicated by the presence of constructions such as ‘taller than’ and ‘very flat’. However, it is not clear that this disanalogy will undermine Cohen’s partners-in-crime response to the objection. For it is not clear that the disanalogies between the context sensitivity of gradable terms and the putative context sensitivity of ‘knows’ makes it more understandable why speakers could be mistaken about the semantic content of their attributions of gradable-adjectives than they could be of their uses of knowledge attributions. Nonetheless, I think Cohen’s partners-in-crime response to the objection is unsatisfactory for two independent reasons.

The first reason is that, it seems to me that if we accuse ordinary speakers of contradicting themselves in cross-context cases where they attribute flatness in one context but deny it in

⁹³ This is effectively the objection pressed by Schiffer (1996).

another context, they will quickly appeal to or easily be led to the purpose relativity of the claims as a way to dispel the appearance of contradiction. However, it seems that if we accuse ordinary speakers of contradicting themselves with their apparently conflicting knowledge claims made in contexts with different epistemic standards, such responses are much less readily available. Cohen (2004) acknowledges this difference and pursues what I take to be a promising explanation—viz. that ‘knows’ unlike ‘flat’ is a *normative* term. He writes, “We find [contextualism] much easier to accept in the case of flatness than knowledge, because ascriptions of flatness do not have the normative force that ascriptions of knowledge/justification do”(2004: 193). However, although I agree that ‘knows’ is a normative term unlike ‘flat’, this does not yet explain why ordinary speakers are more hesitant to admit the context sensitivity of ‘knows’ than of ‘flat’ unless we can say why context sensitivity in normative terms is harder to admit than context sensitivity in nonnormative terms. If Cohen’s suggestion is that it is because we value knowledge, then I want to ask why that makes the crucial difference—after all, pool players value flatness, but they would presumably just as easily admit that ‘flat’ is context sensitive as anyone else. I think we will reach a satisfactory explanation here only by spelling out in some more detail the normative nature of knowledge claims in contrast to the nonnormative nature of flatness claims. In the following chapter, I begin to do just that in my discussion of what ordinary knowledge attributions express and why we use them to do that. However, as far as I can tell, we don’t yet have a plausible explanation of the difference between the robustness of (putative) semantic blindness in the case of ‘flat’ and ‘knows’.

The second reason I am skeptical of Cohen’s partners-in-crime response to the objection is that very few philosophers seriously think that the question of whether some surface is flat generates a perplexing philosophical puzzle, but many philosophers have seriously thought that something like the AFC-paradoxes generates a perplexing philosophical puzzle. This suggests that contextualism requires attributing semantic blindness not just to ordinary speakers but also to philosophers whose job it is to reflect carefully on things like the semantics of ‘knows’. One

might suggest that the disanalogy with ‘flat’ here is due to the fact that philosophers don’t care about flatness as much as they care about knowledge. But precisely because knowledge is one of the traditional subjects of philosophic investigation and philosophers tend to be less semantically blind than ordinary speakers, I would have thought that if ‘knows’ is semantically context sensitive in roughly the same way as ‘flat’ philosophers would have discovered this way out of the AFC-paradoxes a long time ago.

At the beginning of this section, I mentioned the natural worry that the contextualist response to the AFC-paradoxes is ad hoc. I said that the contextualist can overcome this worry only if he can provide reasons independent of the paradoxes for thinking that knowledge claims have the type of context sensitivity that he construes them as having. I think we’ve reached a perspective from which we can see two important things. The contextualist’s posit of an implicit context sensitivity in the truth-conditions of knowledge claims is partially supported by the ordinary use of knowledge claims and thus is not completely ad hoc. Nonetheless, as revealed by the disanalogy between ‘flat’ and ‘knows’, the kind of context sensitivity posited by the contextualist is no ordinary kind of context sensitivity. The sort of implicit context sensitivity posited by the contextualist implicates a deep and hard to explain semantic blindness for anyone who resists contextualism or is perplexed by the AFC-paradoxes. This means that contextualism suffers its own form of the problem of unexplained semantic blindness.

IV. THREE CHALLENGES

In this chapter I have considered the three most prominent families of descriptivist accounts of the meaning of knowledge claims. We found each of these accounts to be lacking in various ways. Specifically, I mentioned three different challenges that, in one way or another, fail to be met by invariantism, subject-sensitive invariantism, and contextualism. These were:

Austin's challenge was to avoid an account of the meaning of knowledge claims which makes ordinary epistemic discourse systematically and widely mistaken. Systematic and widespread mistakes are of course possible, but any account of the meaning of a term which implicates such mistakes owes a very good explanation of them. And skeptical invariantist accounts seem incapable of explaining why we utter so many false knowledge claims. Non-skeptical invariantism, subject-sensitive invariantism and contextualism are better able to meet this challenge. However, because subject-sensitive invariantism has to treat some of our ordinary third-person knowledge claims as erroneous due to projection errors, and because non-skeptical invariantism must treat as erroneous our intuitions that we don't know skeptical hypotheses not to obtain, it seems that contextualism comes out slightly ahead vis-à-vis Austin's challenge.

The *challenge of the paradox* was to provide a plausible response to the AFC-paradoxes. We saw that both the hard-line invariantist responses are unstable. We also saw that the subject-sensitive invariantist has to pursue different strategies for the AFC1- and the AFC3-paradoxes, which, although not a knock-down objection, is a ground for concern. Contextualism seems clearly best equipped to meet this challenge, although we worried that the contextualist response to the paradox is potentially ad hoc.

The *challenge of unexplained semantic blindness* was to avoid an account of the meaning of knowledge claims that posits semantic blindness either to ordinary speakers or to philosophers that cannot be explained. Both skeptical invariantists and non-skeptical invariantists have to treat ordinary speakers as semantically blind. In the former case, the semantic blindness is extreme, for ordinary speakers think many of their knowledge attributions are true, but the skeptical invariantist denies this. In the latter case, the semantic blindness is less extreme, for it only concerns situations where we apply standards for knowledge too high. Here the non-skeptical invariantist argues that our reluctance to attribute knowledge is due to a disconnect

between speaker-meaning and literal meaning. This is not a completely satisfactory answer to this challenge, but it is perhaps better than its competitors. For, as we saw, subject-subject-sensitive invariantists must posit unexplained semantic blindness in their treatment of third-person knowledge claims. And contextualists must posit unexplained semantic blindness in their treatment of ordinary intuitions about the dialectical connections between knowledge claims made in different contexts. It's tough to say which of the positions is best vis-à-vis the challenge of unexplained semantic blindness, but clearly none of them completely meet the challenge.

I don't mean to claim that failure to meet any of these challenges provides a knock-down argument against any of the three descriptivist views I have considered here. Nonetheless, much like the challenges facing prominent descriptivist accounts of ethical claims motivated the search for a better alternative, I think these challenges should leave us unsatisfied enough to consider nondescriptivist alternatives in metaepistemology. Thus, it is to this that I turn in the following chapter.

4 EXPRESSIVIST ACCOUNTS OF EPISTEMIC CLAIMS

In the previous chapter, I considered a recently prominent debate in epistemology about the semantics of knowledge claims. We saw the basic contours of each of the three most prominent positions in this debate—invariantism, subject-sensitive invariantism, and contextualism—and I used various drawbacks of these positions to articulate three complex challenges that any acceptable account of the meaning of knowledge claims should meet. Because none of the three prominent positions met all three challenges, I think we should begin to explore other options.

One way to go forward would be to consider more and more complex and subtle accounts of the fact stated by knowledge claims. However, what I find most interesting about this debate is that it transpires on the assumption that knowledge claims are indeed fact-stating. In effect, this is the metaepistemological correlate to the key assumption of all descriptivist accounts of the meaning of ethical claims. In the dialectics of chapters 1-2, it was the denial of this assumption which opened up conceptual space for nondescriptivist accounts of the meaning of ethical claims. So, in this chapter, I want to consider what happens to the correlative debate about the meaning of knowledge claims if we deny the assumption that they are fact-stating.

My aim, of course, is to show that this opens up room for an account of the meaning of knowledge claims that can make real progress in meeting all three challenges articulated at the end of the previous chapter. To do so, I want to trace out the parallel dialectic in

metaepistemology that we have already considered in detail in metaethics. So, in section I, I shall consider the case for epistemic noncognitivism. This will provide the resources for articulating a fourth challenge to any account of the meaning of knowledge claims, which, in section II, I will address by extending the inferentialist account of ethical claims developed in chapter 2 to the case of epistemic claims. The upshot of my reflections in this chapter will be that there is a way to treat knowledge claims as perfectly truth-evaluable and cognitive while maintaining that they express a sort of practical commitment—roughly speaking, it's a commitment to trust someone with respect to a piece of information. I shall then return to the four challenges and argue that my inferentialist expressivist view of knowledge claims has the resources to meet all four challenges.

I. EPISTEMIC NONCOGNITIVISM

At the level of descriptivism, there is an obvious structural parallel between the landscape of metaethical and metaepistemological positions. Universalism, speaker-relativism, and subject-relativism—in metaethics—map fairly closely onto invariantism, contextualism, and subject-sensitive invariantism—in metaepistemology. Moreover, some of the challenges dogging each pair of parallel accounts are similar. This parallel is not exact, but even the inexact parallel and the thought that ethical and epistemic claims are both normative sorts of claims encourages exploring the possibility of nondescriptivist accounts of epistemic claims parallel to the nondescriptivist accounts of ethical claims explored in previous chapters. I suspect that this has not been done because metaepistemologists have assumed that the only possible nondescriptivist account is an analogue of the traditional and crude emotivist form of noncognitivism defended by Ayer and Stevenson, yet such a view of knowledge claims is patently implausible. But, in my view, emotivism is implausible *both* when applied to epistemological discourse and when applied to ethical discourse, for all of the reasons discussed

in chapter 2. Yet surely the failure of emotivism doesn't undermine all attempts to develop a form of ethical or epistemic noncognitivism. In this section, I want to sketch and defend what I take to be the most plausible version of epistemic noncognitivism. In effect, it will be an extension of Gibbard's quasi-realist version of expressivism.

As we've seen, the most worked out version of ethical noncognitivism is due to Gibbard (1990), who argues that ethical claims mean what they do in virtue of their expressing a state of norm-acceptance. His account of norm-acceptance can be viewed as a simple but powerful modification of the relativist view that the truth conditions of ethical claims are implicitly relative to a set of ethical norms. In one sense, all descriptivists about ethical claims can agree that the sentence,

(1) George's starting the war was immoral,

has truth conditions given by,

(2) George's starting the war was forbidden by ethical norms N.

The universalist will say that there is only one N for all tokens of (1), while the subject-relativist will say that N varies depending on features of George and the speaker-relativist will say that N varies depending on features of who is saying (1). Gibbard's view can be seen as a modification on this speaker-relativist view. He argues that ethical claims $\text{express}_{(p)}$ an interlocking composite mental state called norm-acceptance. This state comprises a descriptive belief whose content is (2) *and* a state of accepting the relevant norms N.

In chapter 3, we saw the advantages and disadvantages of this position in metaethics. What I want to do now is use this as a model for a version of epistemic noncognitivism. We've seen that speaker-relativism is parallel to epistemic contextualism. Contextualism, recall, can be summed up as the claim that sentences of the form

(3) S knows that p,

should be understood as having the truth conditions given by,

(4) S's true belief that *p* meets epistemic standards *e*,

where the value of *e* is determined by features of the speaker's context. Rather than explaining the meaning of instances of (3) in terms of expressing a proposition with relativized truth conditions given by (4), as the contextualist does, I suggest that an *epistemic* noncognitivist could treat instances of (3) as expressions of states of epistemic norm-acceptance. It is epistemic norm-acceptance in that the norms accepted are *epistemic* norms, but norm-acceptance is, again, a complex state of mind consisting of a belief with (4) as its truth conditions and the acceptance of the epistemic norms alluded to therein.

When it comes to (at least some) ethical claims, it is pretty clear what the relevant norm is; they tell us that certain actions are obligated, permitted, or forbidden. But, although it is often said that epistemology is a normative discipline, it is less clear what epistemic norms are. A plausible suggestion here is that the norms relevant to understanding the meaning of knowledge claims are norms which entitle beliefs. Specifically knowledge claims could be understood as expressing our acceptance of particular epistemic norms, which when applied to a particular person's belief entitle or don't entitle the belief. So we might refine (4) with:

(4') S is entitled by norms *e* to her true belief that *p*.

So the proposal is that instances of (3) express a complex state of mind which includes a factual belief with (4') as its truth conditions and the acceptance of the epistemic norms *e* alluded to therein.

Is this norm-expressivist version of epistemic noncognitivism at all plausible? The best way to start to address this question is to see how well the view fares against the three challenges raised in our discussion of descriptivist accounts of knowledge claims. Because epistemic norm-expressivism basically incorporates contextualism into its account of the mental state expressed

by knowledge claims, it can claim many of the advantages of contextualism. But, I think, it also avoids, at least to some extent, the correlated disadvantages.

The primary advantages of contextualism were that it could meet Austin's challenge by underwriting and explaining the apparent cross-context variation among ordinary speakers in willingness to attribute knowledge, and that it could meet the challenge of the paradox generated both of the AFC-paradoxes. At both points, however, the contextualist was forced to posit unexplainable forms of semantic blindness and thereby open himself up doubly to the challenge of semantic blindness. (See page 178f. above for precise articulation of these challenges.)

Regarding Austin's challenge, epistemic noncognitivism can also explain the apparent cross-context variation among ordinary speakers in willingness to attribute knowledge. Recall DeRose's (1992) bank cases, where, in a context where not much hangs on it, someone says that she knows that the bank is open on Saturday, but, in another context where a lot hangs on it but she has the exact same evidence, she denies this knowledge. A proponent of the sort of epistemic noncognitivism that I have been sketching would explain this variation much like the contextualist explains it—viz. in terms of a variation of the *e* in the content of the belief, which is part of the state of norm-acceptance. However, the epistemic noncognitivist will attribute this phenomenon not to cross-context variations in the operative epistemic standards but to cross-context differences in which epistemic norms are accepted by those attributing and denying knowledge. This can, in turn, be explained in terms of contextual variations in our reasons for evaluating whether or not a particular believer is entitled to his beliefs.

Contextualism achieves its answer to Austin's challenge at the cost of positing a troubling form of semantic blindness. This is because, as we saw in the previous chapter, contextualism has to disrespect ordinary intuitions about the dialectical connections between some cross-context knowledge claims (and their metalinguistic variants). But the epistemic noncognitivist position that I have been sketching has a way out of this problem. The basic idea is that by

maintaining that knowledge claims express states of norm-acceptance rather than relativized descriptive beliefs, the epistemic noncognitivist like the ethical noncognitivist gains a second axis of possible opposition or agreement. That is, he can account for the intuition of cross-context dialectical connections between knowledge claims by pointing out that, on his view, knowledge claims can express pragmatically opposed or concurring states of norm-acceptance. Finding logically contradictory or identical descriptive beliefs is not the only way to account for dialectical connections.

More intuitively, it is easy to see how a speaker's willingness to treat someone as entitled to her belief that *p*—i.e. acceptance of particular epistemic norms—might vary with the speaker's changing circumstances. If I am just idly speculating about whether *S* knows that *p* and am not especially worried about possible defeaters, I will be perfectly willing to treat someone as entitled to her belief even if it is based on good but perhaps not terrific evidence. However, if my circumstances change and I need to rely on *S* as an informant about whether *p*, I may become worried about possible defeaters; then, if I suspect that the person doesn't have terrific evidence for her belief, I will be less willing to treat her as entitled to the belief. In this way, the norm-expressivist can explain the intuitive variations in the attributability of knowledge across conversational contexts without positing the sort of semantic blindness about dialectical connections that the contextualist was forced to posit.

Of course, different epistemic norms might entitle beliefs for importantly different reasons. For example, we could imagine reliabilist norms which entitle beliefs formed by a reliable belief forming process and reason-based norms which entitle beliefs which are based on sufficient reasons. So, in some ways, this proposal is neutral with respect to traditional epistemological debates about the nature of knowledge. In fact, I suspect that part of the reason that such traditional debates have seemed so intractable is that metaepistemologists have illicitly assumed that knowledge claims are descriptive and thus taken it for granted that knowledge claims have to be understood in terms of the attribution of some constellation of factual properties, which in

turn encourages the thought that there is some fact of the matter about whether reliabilists or internalists are right about the nature of the worldly relation of knowledge. However, if we give up this assumption in epistemology, as the expressivist tradition has urged us to do in ethics, we gain the resources to explain why the differences seem intractable, while neither resolving them nor dismissing their importance.⁹⁴

Regarding the challenge of the paradox, the epistemic noncognitivist can mimic the contextualist resolution of the paradoxes. The core idea behind the contextualist's strategy for resolving the AFC paradoxes was to treat the conclusion of the skeptical AFC arguments as true *in the context in which the argument is presented*; but this is consistent with the view that many ordinary knowledge attributions are true *in the more quotidian contexts in which they are made*. Whatever plausibility this response to the AFC-paradoxes gains for contextualism can be just as easily captured by the epistemic noncognitivist position I have been sketching. The strategy will be to transmute the contextualist's notion of varying contextual standards for knowledge into the notion of varying states of epistemic norm-acceptance. That is to say that the epistemic noncognitivist will argue that the AFC arguments have the pull that they do insofar as they get us to accept more demanding epistemic norms than we typically do in ordinary contexts. But the resistance to their conclusion is due to the fact that we harbor acceptance of the less demanding standards.

Contextualism achieves its answer to the challenge of the paradox at the cost of positing semantic blindness to anyone who has been troubled by the AFC-paradoxes, which generates the worry that it is an ad hoc response to the paradoxes. Importantly, however, the epistemic noncognitivist, unlike the contextualist, treats the skeptical conclusion as *in conflict* with ordinary knowledge claims; and this explains why, in a sense, it is right to be troubled by the

⁹⁴ For an independent argument that at least some epistemological discourse is partially nonfactual, which is motivated by similar observations about the intractability of certain sorts of epistemological debates, see Field (1998).

paradox. To be sure, for the norm-expressivist, the relevant conflict is not (typically) a conflict between logically contradictory descriptive beliefs but rather a disagreement in which norms are accepted. A further advantage of this response to the AFC-paradoxes is that it can explain why some people find the paradoxes more troubling than others. Insofar as one is either a convinced skeptic or a convinced anti-skeptic, the contingencies of, respectively, ordinary epistemic discourse and skeptical arguments will not sway one's acceptance of particular epistemic norms. But if one's acceptance of particular epistemic norms is more flexible, then one will more easily be drawn into a situation where it seems as if some knowledge attribution is both true and false. This allows for a metaepistemological theory that is importantly neutral with regard to certain normative epistemological debates. For by making the meaning of knowledge claims depend on the states of norm-acceptance expressed by whoever makes them, the epistemic noncognitivist is not forced to take sides on who is right in a debate between skeptics, anti-skeptics, and those who cannot make up their minds.

So far, epistemic norm-expressivism appears to answer Austin's challenge and the challenge of the paradox just as well as contextualism, and, since contextualism answered these challenges better than any of the various invariantist positions, this suggests that norm-expressivism is better than them with respect to the first two challenges. Moreover, epistemic norm-expressivism appears to meet the challenge of semantic blindness better than contextualism, and so it does not have the central drawback of that view.

We should recognize, however, that the norm-expressivist strategy for meeting the challenge of semantic blindness is not completely satisfactory. In effect, the epistemic norm-expressivist co-opts the contextualist's answer to Austin's challenge and the challenge of the paradox, but then he adds an extra noncognitive element—norm-acceptance—in order to avoid positing semantic blindness. I think this is a better package than contextualism, but an opponent could complain that this is an ad hoc maneuver. For it appears to generalize and provide a formula for treating any cross-context disagreement as nonfactual.

Take, for instance, the disagreement between mainstream evolutionary biologists and proponents of the intelligent design theory of biological development. This seems to be a disagreement about a matter of fact—viz. what the mechanism that regulate phylogenetic biological development is. Mainstream biologists say that it is random mutation controlled by pressures towards reproductive fitness. Proponents of intelligent design say that it is mutation directed by the design of an intelligent being. What, then, would we think of a would-be *biological* norm-expressivist who argues that this disagreement is not a disagreement about a matter of fact rather a disagreement in which norms are accepted by the two parties? Surely the two parties do accept different norms of scientific investigation, but that does not seem to be enough to capture the nature of their disagreement. For it is most plausible to think of the norms of scientific investigation as norms governing the production of a univocal theories about aspects of *the* (one and only) natural world. So a disagreement in scientific theory may indeed result from a disagreement in scientific norm-acceptance; however, given the nature of scientific norms, it is implausible to think of this disagreement as nonfactual or a matter of a difference in pragmatically opposed noncognitive states.

In order to overcome this worry, I think the epistemic norm-expressivist would need to provide a plausible account of the social role of accepting particular epistemic norms and expressing these states of epistemic norm-acceptance; and this account would need to construe these norms as organized around something other than the goal of producing a univocal theory of the epistemic aspects of *the* natural world. Of course, Gibbard (1990) provides a sustained attempt to do this for the ethical case. Maybe some of his account could be carried over to the epistemic case, but no one has yet tried.

Perhaps this is because there is another serious worry about epistemic noncognitivism lurking in the wings. The view treats knowledge claims as noncognitive, but knowledge claims are typically taken to be paradigmatically cognitive. Traditionally, thinking of a claim as cognitive has meant thinking of it as truth-evaluable and expressive of beliefs. And it seems that

all parties to this debate should admit that knowledge claims are cognitive, in this sense. For the AFC-paradoxes wouldn't even appear to be paradoxes if we didn't take the triad of individually plausible yet mutually inconsistent claims to be straightforwardly truth-evaluable and expressive of beliefs. After all, what it is for a set of claims to be *inconsistent* is for it to be impossible that they are all *true*. And the reason these are paradoxes is that we take ourselves to have good inter-subjectively valid reasons for *believing* that each of the claims is *true*. Moreover, ordinary speakers will metalinguistically ascribe truth to knowledge claims, embed knowledge claims in belief contexts, and reason with their knowledge claims by, e.g., embedding them in conditionals and running modus ponens arguments. If the epistemic noncognitivist denies all of these apparently cognitivist features of ordinary and philosophical discourse about knowledge, he is surely involved in a drastic posit of semantic blindness, which outweighs any other advantages of his account.

Notice, however, that this worry turns on features of epistemic discourse parallel to the apparently cognitivist features of ordinary and philosophical discourse about *ethics* that we saw causing problems for traditional forms of ethical noncognitivism in chapter 2. There, we also saw that, under these pressures, traditional noncognitivism has evolved into contemporary “quasi-realism”. I articulated this evolution by distinguishing between two senses of ‘express’. In the (classically) semantic sense, declarative sentences are typically thought to express_(s) propositions, while, in the (initially) pragmatic sense, utterances of declarative sentences are typically thought to express_(p) the speaker's beliefs. We saw that traditional noncognitivism comprised a negative thesis about ethical claims at the level of both expression_(s) and expression_(p). That is, noncognitivists hold that ethical sentences do not express_(s) a proposition, and ethical utterances do not express_(p) a belief. Distinguishing these two claims, quasi-realism can, then, be seen as a denial of the former and a reinterpretation of the latter.

As we saw, quasi-realists deny the negative claim about what ethical sentences $\text{express}_{(s)}$ by adopting a minimalist stance towards notions such as truth and propositionality. So, they have no problem recognizing that, in a minimal sense, ethical sentences $\text{express}_{(s)}$ propositions and are thus perfectly truth-evaluable like any other declarative sentence. However, we also saw that many of the same reasons that support the quasi-realist's minimalist stance about truth and propositionality bled into accounting for the psychological state $\text{expressed}_{(p)}$ by ethical claims. The crucial point about this was that it looks as if quasi-realists must admit that, in a sense, ethical claims $\text{express}_{(p)}$ beliefs. Even so, however, they maintain that they have an interesting distinct account of the state $\text{expressed}_{(p)}$. They'll say that, whether or not we call it a belief, it has a distinctive evaluative noncognitive element.

I see no reason to think that a similar quasi-realist strategy is not available to the *epistemic* noncognitivist, which would quell the crude although powerful worry that any form of epistemic noncognitivism treats what is obviously cognitive (knowledge claims) as noncognitive. However, going this route incurs the problems of quasi-realism more generally. Specifically, as we saw in chapter 2, quasi-realists face the challenge of creeping minimalism and valid reasoning. This was the problem of nonarbitrarily stopping the minimalist deflation of the ontological significance of terms such as 'truth', 'proposition', 'belief', 'representation', etc. Without a definite stopping point, quasi-realism cannot clearly distinguish itself from its cognitivist competitors, and so threatens to undermine the whole debate between descriptivism and nondescriptivism. Moreover, as in the ethical case, it is unclear that epistemic quasi-realism would be able to make sense of ordinary reasoning with knowledge attributions. Why, for instance, are the AFC arguments valid? This is the Frege-Geach worry transmuted into the epistemological arena; and it is not any clearer here that quasi-realists have a satisfactory answer in the offing.

In direct analogue to the dialectic of chapter 2, I'd like to take the worries expressed in the preceding paragraphs as a further challenge to an acceptable account of the meaning of

knowledge claims. Since it's the same challenge transferred from the metaethical to metaepistemic arena, I shall give it the same name—viz. *the challenge of creeping minimalism and valid reasoning*. It threatens only nondescriptivist accounts that, in order to respect the surface phenomena of ordinary discourse, adopt a minimalist stance towards terms like 'truth' and 'belief'.

The inferentialist expressivist account of ethical claims developed in chapter 4 offers a plausible nondescriptivist account that overcomes the challenge of creeping minimalism and valid reasoning. My leading thought going forward is to develop a parallel inferentialist expressivist account of knowledge claims, in order to see whether it can meet the four challenges to an acceptable metaepistemological account that have emerged from our discussion thus far.

II. EPISTEMIC INFERENTIALIST EXPRESSIVISM

In chapter 2, I pointed out that quasi-realist expressivists tacitly presuppose a certain kind of semantic explanation. They operate under the assumption that the meaning of a claim is to be given in terms of what mental state it conveys. This "ideational" explanation of semantic content is, of course, one of the classic traditions in the philosophy of language. However, it faces a number of well-known problems. Moreover, in the debate about the viability of a nondescriptivist account of knowledge claims, presupposing this account of meaning seems to be precisely the problem with quasi-realist noncognitivist accounts. For either the semantically relevant state expressed by a knowledge claim is a sentiment or emotion, in which case it is unclear why it can serve as the sort of thing that we can reason to and from and why it gets treated as a belief in ordinary discourse, or the state is a genuine belief, in which case it is unclear what the difference is between quasi-realism and straightforward descriptivism. In short, the assumption of an ideational theory of meaning forces epistemic quasi-realism to be

either too much like traditional noncognitivism to be plausible or too much like descriptivism to represent a distinctive nondescriptivist alternative.

These are not decisive reasons to abandon the ideational theory of meaning. However, if some other theory can bring new resources to bear on the standoff between descriptivists and quasi-realists, then that will provide indirect support for that theory and, more importantly, help us to see a way towards a more plausible account of the meaning of knowledge claims. The theory that was able to an analogous job in the case of ethical claims was the sort of semantic inferentialism propounded by Frege (1879), Sellars (1948/1980, 1968), Rosenberg (1974), and Brandom (1994). The core idea is that the linguistic meaning of a claim is to be given in terms of the sorts of inferences it can figure in, both as a premise and as a conclusion.

Recalling the distinction between $\text{expressing}_{(s)}$ and $\text{expressing}_{(p)}$, semantic inferentialism can be brought to bear on the standoff between descriptivists and quasi-realists by noticing that in addition to talking of claims $\text{expressing}_{(s)}$ propositions made true or false by the facts, as the descriptivist does, and of claims $\text{expressing}_{(p)}$ a special kind of mental states, as the quasi-realist does, we also talk of claims expressing *commitments*. (Depending on how we individuate expression-relations, this could be seen as a third kind of expression or a non-psychologistic variant on $\text{expression}_{(p)}$.) Now there are a lot of ways to understand this locution; however, a plausible and useful way in this context is as follows: We can refer to the commitment expressed by various claims in a quasi-disquotational way. So, for instance, I'll say that the claim "It's raining" expresses a commitment to it's raining, and the claim, "Jerry knows what time it is" expresses a commitment to Jerry's knowing what time it is. But then, we can explain the nature of these commitments in terms of the inferential roles of the claims that express them. An inferential-role of a claim, recall, is a function of the upstream and downstream inferential potential of the claim (see page 114f. for more description). In this framework, we can then say that the semantically relevant sort of commitment expressed by a claim is this inferentially

articulated commitment. If two claims have the same upstream and downstream inferential potential, then they express the same commitment, otherwise they express different commitments.

This strategy provides completely general account of the semantically relevant commitments expressed by both ostensibly descriptive and nondescriptive claims. However, within this inferential framework, we can still achieve the sort of contrast important to expressivists. The key is to deploy the distinction between theoretical and practical inferences. Some inferences move from theoretical premises to a theoretical conclusion, while other inferences move from a combination of theoretical and practical premises to a practical conclusion. The commitments that can serve as practical premises and practical conclusions thereby have a distinctive practical element, which, in the present terminology, we can capture by saying that they are practical commitments. This distinguishes them and the claims that express them from theoretical commitments and the claims that express them.

In the dialectics of chapter 2, this distinction provided an attractive way to generate a new form of ethical expressivism. The leading idea was to treat ethical claims as expressing practical commitments, which distinguished them ontologically and epistemologically from descriptive claims which were seen as expressing theoretical commitments. Admittedly, it is far less obvious that knowledge claims should be treated as expressing practical rather than theoretical commitments; however, two considerations speak in its favor. First, beside ethical claims, knowledge claims are paradigms of normative claims. So if treating ethical claims as expressions of practical commitments provides a way in metaethics to capture and begin to understand their distinctive normativity—which I think it does—then a similar strategy should appear attractive in metaepistemology. Second, the standoff between descriptivists (of various sorts) and quasi-realists in *metaepistemology* discussed in the previous chapter and the first half of this chapter are similar enough to the standoff between descriptivists (of various sorts) and quasi-realists in

metaethics that it would be highly surprising if a third kind of account available in the metaethical debate were not also available in the metaepistemological debate.

The reason it is less obvious that knowledge claims should be treated as expressing practical rather than theoretical commitments is that it is less obvious what sorts of actions can be inferentially supported by a knowledge claim in the way distinctive of a practical premise. So what I want to do next is to work towards an account of the type of action I think is inferentially linked to knowledge claims by discussing three recent proposals for understanding the essentially practical nature of knowledge claims. The first of these comes from Edward Craig.

In reaction to the philosophical annals full of more and more complex proposed analyses of the concept of knowledge, Craig (1990) suggests a different angle on the traditional epistemological project of understanding knowledge. He proposes to ask what the point of the concept of knowledge is. He recognizes that the point of some of our concepts—e.g. the concept of water—is to pick out or track certain natural stuffs; and, if the concept of knowledge were like that, then no more understanding of knowledge would be gained by learning about the point of the concept of knowledge than is gained of water by learning that the point of the concept of water is to pick out or track water. However, Craig thinks the point of the concept of knowledge is importantly different.

He proposes to think about knowledge much like Hobbes thought about justice. That is to say, roughly, that the concept of knowledge can be understood as the solution to a certain sort of social coordination problem that we face as human beings. Craig calls this a ‘state-of-nature’ explanation of a concept. He writes,

We are attempting a ‘state of nature’ explanation of a number of facts of conceptual or linguistic practice. Such explanations work by identifying certain human needs and arguing that the practices are a necessary (or at least a highly appropriate) response to them; they will therefore be at their strongest when the human needs from which they start are the most practical, hence the most undeniable ones.(Ibid.: 89)

In the case of justice, Hobbes famously argued that, given natural human psychology, we need a way, as Craig puts it, “to regulate power-relationships between individuals...to avoid the otherwise inevitable ‘war of every man against every man’ and create stable circumstances in which humans, and human society, may flourish”(Ibid.: 9). And, for Hobbes, justice then is whatever social institutions satisfy this need.

Similarly, Craig’s hypothesis is that knowledge is the solution to a different social coordination problem that we face as human beings. He describes the problem as follows:

Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths. They have ‘on board’ sources, eyes and ears, powers of reasoning, which give them a primary stock of beliefs. It will be highly advantageous to them if they can also tap the primary stocks of their fellows...that is to say, if they act as informants for each other...So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use.(Ibid.: 11)

Hobbes argued it would be highly advantageous for humans to have the social institutions which adjudicate between the incompatible interests of all members of a community. Likewise, it seems right, as Craig suggests, that it would be highly advantageous for humans to be able to rely on their fellows as informants of pieces of information not immediately available to the ‘on board’ information sources. Thus, as Hobbes argued that justice just is the solution to the problem of coordinating in community incompatible interests, Craig argues that knowledge just is the solution to the problem of keeping track in community of who are good informants. He writes, “To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information”(Ibid.: 11).

In the project of identifying the practical point of knowledge claims by isolating the types of actions inferentially connected to knowledge claims, Craig’s account provides a helpful initial proposal. The idea would be that knowledge claims are the sorts of claims it takes to inferentially support the action of treating someone as a good informant. This may involve various things, such as relying on S’s testimony as to whether p, citing S as the source of one’s information

about *p*, etc. And it would seem that, in the terminology of chapter 2, knowledge claims should come out as *collective* practical commitments rather than *individual* practical commitments—i.e. their content should be modeled in terms of ‘we-shall’ commitments rather than ‘I-shall’ commitments. This is because they justify the action of treating someone as a good informant not merely for the particular inquirer *A* who attributes knowledge to *S* but rather for anyone in *A*’s position of looking for a good informant about *p*. This suggests that the content of a knowledge claim of the form,

(3) *S* knows that *p*,

is something like,

(3’) We shall treat *S* as a good informant as to whether *p*.

Now, however we come down on the details about whether the practical commitment expressed by knowledge claims is correctly modeled in terms of ‘we-shall’ or ‘I-shall’ commitments, there is a flaw in this line of thought that is important to recognize. As long as *S* is someone other than the speaker, it doesn’t seem to make sense to view instances of (3) as “flagging” a good informant as to whether *p*. For anyone willing to assert or accept (3) *already* takes *p* to be true and so won’t need (or need to flag) an informant about the truth value of *p*. I think these considerations reveal that Craig’s account is better suited for claims of the form,

(5) *S* knows *whether* *p*.

To be sure, instances of (5) are one species of knowledge claims that any satisfactorily complete metaepistemological account will have to have a place for. But I am skeptical of the prospects of elucidating instances of (3) in terms of Craig’s ‘state-of-nature’ account that works most clearly only for instances of (5). It may work in the first-personal case where by saying “I know that *p*” one is simultaneously flagging oneself as an approved source of information and transmitting the information one is approved to transmit. However, that is a special case and

doesn't generalize to an account of (3) in terms of (5). This is not to say that I think Craig is wrong to stress that in thinking about the real human practice of attributing knowledge, "we must never forget that the inquirer's situation is a practical one..."(Ibid.). It's just that I think inquirers face a variety of practical problems, and so even if sentences of the form (5) serve the practical purpose of, roughly, flagging approved sources of information, that doesn't imply that all sentences of form (3) serve the same purpose.

So, in the project of understanding the practical point of knowledge claims, I think we should provisionally endorse Craig's account as applied to (5). That is to say that saying that someone knows whether p can inferentially support the action of treating them as a good informant about p. But in order to reach a satisfactory account of knowledge claims in general, I think we need to look elsewhere for the practical point of claims of the form (3). The next place I propose to look is to some of Jay Rosenberg's recent work.

In his 2002, Rosenberg is concerned to argue against what, in some (wide) philosophical circles, appears to be the received view about knowledge. This is, roughly, the view that knowledge is a "determinate state of a person—a state which results from being related in certain determinate ways to the world"(2002: 196). He takes this view to result from, *inter alia*, the Cartesian beginnings of modern epistemology, which conceive of epistemological investigations about whether we know what as detached from the specific situated epistemic practice of various inquiries. In the Cartesian tradition, "[t]he task of a philosophical account of knowledge...[is] to fill in the details concerning the conditions for being in that state [i.e. knowledge], and thereby, *inter alia*, to arrive at an assessment regarding one's chances, *anyone's* chances, of realizing them"(Ibid.: 196). Viewing the central task of epistemology this way leads, in Rosenberg's view, to a number of problems—not least of which is the nagging problematic of Cartesian skepticism.

Here, I shall not discuss Rosenberg's arguments against various moments and twists in this Cartesian tradition. What is more important for my purposes is the positive account of knowledge that he proposes as an alternative to the family of views which emerges from Cartesian tradition. There are three key components: his account is resolutely anti-skeptical, thoroughgoingly fallibilist, and perspectivalist (2002: 2). Here I shall focus mostly on the latter component. Rosenberg writes,

The "*perspectivalist*" view here advocated, whose origins line in Peircean pragmatism, is predicated upon explicit recognition of our inescapable epistemic *situatedness*. Unqualified attributions of knowledge are always made from our own epistemic perspective, with reference to a particular context of enquiry, and *enquiry* itself is correlatively understood as always addressed to determinate questions, properly raised only within the framework of a set of background beliefs, themselves not then and there in question, which render appropriate the application of particular investigative procedures, techniques, and norms of epistemic conduct.(Ibid.)

The sense in which Rosenberg sees us as epistemically situated comes out most clearly in his discussion of Stroud's (1984) critique of Moore's anti-skeptical arguments. Stroud's key complaint against these arguments is that Moore fails to recognize the possibility of an "all at once" skeptical challenge resulting from questioning the very possibility of knowledge from a (meditative) perspective detached from the exigencies of everyday life or normal inquiry. Stroud writes,

[What] Moore ignores is the possibility of a detached philosophical assessment of the assertions and knowledge-claims we make in everyday life. The conclusion of that assessment is that none of those confident assertions are strictly speaking instances of knowledge, and so none of our knowledge-claims about the world around us are strictly speaking true. There is a direct conflict between the philosophical conclusion and our claims in everyday life.(1983: 176-177)

According to Rosenberg, however, Moore doesn't *ignore* this possibility of a detached assessment—he *denies* it on the grounds that "what the sceptic's ostensible 'detached assessment' is detached *from* turn out to be just the considerations which give epistemic assessment its point in the first place"(2002: 174).

Clearly this response to Stroud rides on what considerations indeed give epistemic assessment its point, and we've already (with Craig) begun to see a conception of the point of knowledge claims that makes the skeptic's detached perspective appear less than completely coherent. For, if knowledge claims are assessments of good and bad informants in our own pursuit of new information, then knowledge attributions would have no point when we detach from this pursuit.⁹⁵

Rosenberg pursues a similar thought. To this end, he mobilizes the idea of a *community of inquiry*, writing, "At any given time such a community will have some *open questions* so to speak 'on its agenda'. That is, there will be propositions whose truth or falsehood is an object of more-or-less active and ongoing enquiry"(Ibid.: 187). Also, however, "at any given time there will be propositions whose status for a community of enquiry at that time is, at least provisionally, effectively settled—either as expressing truths that are available for use as premises and data or as expressing falsehoods whose contradictories can be so used"(Ibid.: 187-188). This is because such a community couldn't get anywhere in their inquiry into the truth or falsity of the open questions unless they rely on some set of answers to questions that they take to be closed. Moreover, such a community couldn't get anywhere in their inquiry if they didn't have what Rosenberg calls an *epistemics*—i.e. "a family of relevant and acceptable *procedures* of enquiry; that is procedures for transforming propositions from open questions into propositions with settled status"(Ibid.).

In this account of situated inquiry, the point of epistemic assessments is to keep track of which questions are open and which questions are closed in the ongoing inquiry. Some questions that were open will become closed, and, if the inquiry changes course, some questions that were closed may be opened (as long as enough other questions remain closed to continue

⁹⁵ Of course, in a particular localized case, we may think that no one around us is a good informant for our current pursuit of some piece of information. However, it seems that we detach completely from this pursuit as soon as we think that no one *could* be a good informant.

inquiry). Rosenberg writes, “since propositions can change status in this way...it will be useful for such a community to have a means of *keeping the books*. It will be useful, that is, for the community to have available a way of indicating, from its collective epistemic perspective, the status of individual propositions relative to the ongoing enquiry in light of the pertinent epistemics”(Ibid.).

From the point of view of our investigation into the practical point of making knowledge claims of the form

(3) S knows that p,

Rosenberg’s account suggests a proposal helpfully different from Craig’s. The idea would be that knowledge claims are inferentially tied to the action of treating someone as having legitimately closed a question in an ongoing inquiry of which we are, at least implicitly, a part. This presumably involves various things, such as citing the answer to the question as evidence for closing other questions, not challenging S’s belief about the answer of the question, not investigating further into the answer to the closed question, disregarding skeptical challenges about the answer to the question, etc. Moreover, again, the type of practical commitment apparently expressed by a knowledge claim would be of the collective rather than individual variety. More specifically, if we are modeling the commitment in terms of ‘we-shall’s, we’d get something like,

(3") We shall treat p as legitimately closed for S.

Again, however, whether or not we accept this precise modeling of the practical commitment expressed by knowledge claims, there is a flaw in this line of thought that is important to recognize. Explaining (3) in terms of treating someone as having legitimately closed a question in an ongoing inquiry of which we are, at least implicitly, a part works well for the species of (3) where the relevant ‘S’ is a member of our present community of inquiry. That is, Rosenberg’s account can explain the point of claims of the form

(6) It is known (by us, me, or my fellow inquirer) that p.

These might be modeled by something like,

(6'') We shall treat S as having done everything necessary in our present inquiry to legitimately hold it as the case that p,

where the 'we' refers to members of the relevant community of inquiry. However, surely this is not the only species of (3). For even if, as Rosenberg claims,

The basic role of appeals to the concept of knowledge is...to keep the books on epistemic entitlements and responsibilities *vis-à-vis* an enquiry aimed at closing *particular* open questions by satisfying the requirements of a *specific* epistemics against a background of shared beliefs *not* then and there in question.(Ibid.: 197)

we seem to “keep the books” on a whole range of people, including people who are not plausibly considered part of our present “community of inquiry”. Does George Bush know that there are no weapons of mass destruction in Baghdad? Well, if he’s trying to justify the war, we might say that he does know this and so cannot use the original justification; however, if he is trying to plan a visit to Baghdad, we might say that he doesn’t know this and should therefore take extra precautions. This seems like perfectly ordinary epistemic question, yet the answer to it could vary depending on the inquiry in which the relevant epistemic agents—*rather than we*--are engaged. Moreover, George Bush may not share an epistemics with us nor consider the questions presently open in our community of inquiry (say—whether George Bush should be impeached) to be open questions.

These reflections appear to show that, as far as it has been elucidated, Rosenberg’s account works nicely to explain the practical point of claims of the form (6). The idea is that claims of this form are inferentially connected to the action of treating the question as to whether p as closed in the ongoing inquiry. However, I think this account (like Craig’s account) fails to

capture the practical point of knowledge claims in its full generality. So I'd like to look at one more proposal from Timothy Williamson.

Williamson has written much about knowledge in the past decade. Most notably his 2000 defends the view that knowledge is a simple mental state, unanalyzable into component parts. However, in two more recent articles (2005a and 2005b), where he takes on contextualist accounts of knowledge attributions, he has sketched a view that links the truth of knowledge attributions closely to practical reasoning. It is reflection on this view that I think will help to fill in our understanding of the practical commitment expressed by knowledge claims.

As Williamson sees things, epistemic contextualism generates a counterintuitive result on par with a counterintuitive result generated by metaethical speaker-relativism. I discussed metaethical speaker-relativism in more detail in chapter 1, but for present purposes it suffices to recall that the speaker-relativist's view is that a claim, such as,

(7) It would be wrong for Clare to resign from her job,

has truth conditions given by,

(8) Clare's resigning from her job is forbidden by norms N,

where the value of N is determined by features of the speaker of (7) in such a way that different speakers or the same speaker in different contexts can truly utter (7) and not-(7). Williamson argues that this view is deficient because it privileges the speaker's context over the agent's context in an obviously practical matter; and in practical matters the context of the agent at the time of action must be privileged in order to make sense of the relevance of certain considerations to practical decision making.

To make this point, Williamson considers Clare herself thinking (7) when certain considerations are salient and, then, when other considerations are salient, changing her mind and thinking not-(7). Because Clare is committed to not doing the wrong thing, she is torn. In

the relativist's voice, Williamson then presents the following counsel to Clare's vexed deliberations:

Do not worry. You are mistaken in supposing that there is a disagreement between what you think when you think 'It would be wrong for me not to resign' and what you think when you think 'It would not be wrong for me not to resign'. Both thoughts are true. The sentence 'It would be wrong for me not to resign' expresses different propositions as uttered in different contexts.(2005b:103)

But Williamson objects: "This...resolution of Clare's problem is liable to strike us as glib and shallow. When the moment for action comes, Clare must either resign or not resign...She cannot have it both ways"(Ibid.).

Here, because of Clare's (practical) commitment to not doing the wrong thing, it seems that the only context that matters to the sense of 'wrong' is the context of Clare's deciding and then acting. Any other context is either irrelevant or parasitic on this context. As Williamson puts the point, "The speaker's context lacks autonomy...in setting the content of the practical question. In some sense 'Would it be wrong?' is above all a question *for* the agent, in the context of agency (which is not to deny that the agent may fail to raise it); it is a question for others, or for the agent in other contexts, *because* it is a question for the agent in the context of agency"(Ibid.: 105).

In Williamson's view, the situation is parallel when it comes to 'knows'. He writes, "In some sense 'Does S know at *t* that P?' is above all a question for the agent S, at time *t* (which is not to deny that S may fail to ask the question); it is a question for others, or for S at other times, *because* it is a question for S at *t*"(Ibid.: 107). He offers two different arguments for this view. In 2005b, he suggests the following principle:

- (9) If one knows that P, then one can hardly be wrong to believe that P; conversely, given that one does not know that P, it arguably is wrong to believe that P.(Ibid.: 108)

Then he constructs a case in which Mary believes on the basis of memory that she had her purse yesterday morning. Does she know that she had her purse yesterday morning? For the contextualist, whether she knows will depend on—in addition to the truth-value of Mary’s belief and the strength of the evidence on which it is based—the context in which the question is asked. But, given (9), contextualism then implies that whether it is wrong for Mary to believe that she had her purse can vary in very much the same way that the ethical speaker-relativist is committed to saying that whether it is wrong for Clare to resign her job can vary. But Williamson claims that, in the case of Mary, “[t]he primary question seems to be ‘Is it wrong for me to believe that I had my purse yesterday?’ as uttered by Mary.

There is, however, an important disanalogy between the two cases, which Williamson fails to note. In the former case, Clare is trying to decide what to do. That she is committed to not doing what is wrong makes her question of whether resigning her job is wrong a distinctively practical question. In the latter case, Mary *already* believes that she had her purse. She is not in what Williamson calls a “context of agency”. Indeed, it is far from clear that we are ever in a context of agency vis-à-vis our own present beliefs—for one cannot believe at will. But, even if we bracket worries about doxastic volunteerism and assume that one could face the practical question of whether to believe P and be committed to not believing wrongly, in such a context one would not yet believe that P and so *ipso facto* not know that P. Given (9), this implies that one would be wrong to believe that P, whenever one does not already believe P. But clearly that is absurd. For, in many cases, were one to come to believe that P, then one would know that P.

I take this to be a *reductio* of (9), but I don’t think it undermines an idea close to Williamson’s—viz., that when S believes that P, if S knows that P, then S’s belief that P is somehow appropriate, while, if S doesn’t know that P, then S’s belief that P is somehow inappropriate. Of course, making this idea more specific requires fleshing out the relevant sense in which S’s belief might be appropriate or inappropriate. And there is a clear and natural sense

in which this idea is misguided. One's belief might be appropriate if it is *epistemically justified*, but a false belief can be epistemically justified,⁹⁶ and such beliefs clearly do not constitute knowledge.

So we need a different sense of appropriateness of beliefs if we are to make sense of this proposed connection between knowledge and the appropriateness of beliefs. In his 2005a, Williamson himself suggests that the connection must be to the agent's practical reasoning. He writes,

There is an intuitively attractive link between knowledge and practical reasoning. If it is not true for me to say 'I know that my door is locked', something is wrong with practical reasoning in which I rely on the premise that my door is locked, for in some sense I am not entitled to that premise, however much it seems to me that I am.(2005a: 227)

He recognizes that even when the belief is false one may have good evidence or justification for thinking that it is true, but he insists that this doesn't mean that one is entitled to relying on the belief in practical reasoning. For being justified in a belief may provide a good *excuse* for relying on it without providing entitlement. Williamson offers the following comparison: being justified in thinking that one can take books out of the library may provide a good excuse for doing so even if it turns out (perhaps due to a recent change in the rules) one is not entitled to do so.

So Williamson's new proposal is to understand the relevant sense in which beliefs can be appropriate or inappropriate in terms of whether one is entitled to rely on them in practical reasoning. This generates the following principle:

(10) S knows that-*q* iff S is entitled to rely on her belief that *q* in her practical reasoning.⁹⁷

⁹⁶ At least this is true on most conceptions of epistemic justification. For a contrary view, see Sutton (2005).

⁹⁷ This is a variant of the principle Williamson names 'KPR+'—see Ibid.: 231. KPR+ is a generalization from his principle KPR: "A first-person present-tense ascription of 'know' with respect to a proposition is true in a context iff that proposition is an appropriate premise for practical reasoning in that context"(Ibid.: 227). In Williamson's discussion, the key difference between KPR and KPR+, besides the generality of the latter, is that KPR is already relativized to a context. However, as he intends it, 'know' in KPR+ should be interpreted invariantly.

Williamson presents this not as an analysis of the concept of knowledge but, nonetheless, as necessary and sufficient condition for knowledge.

It is not immediately obvious, though, how Williamson is conceiving of relying on a premise in practical reasoning. For on the classical syllogistic conception of practical reasoning, a typical piece of practical reasoning moves from a major premise expressing an intention to achieve some end,

(P1) I shall achieve end E

and a minor premise expressing a belief about means⁹⁸,

(P2) Doing M is the best way to achieve E,

to reach a conclusion, which is the expression of a more proximate intention (or, on some accounts, even the action itself):

(C) Thus, I shall do M.

It is controversial whether there are other valid forms of practical reasoning, but everyone should agree that the classical practical syllogism formalizes a clear *paradigm* of valid practical reasoning.⁹⁹

Within the practical syllogism, there are only two kinds of premises that it might be appropriate to rely on. First, there is the major premise expressing an intention to achieve some end. Second, there is the minor premise expressing a belief connecting means to ends. It is debatable whether the major premise of the practical syllogism is even the sort of thing that can constitute knowledge. It's important to note that it expresses not a predictive belief that one *will* achieve some end and not a self-ascriptive belief that one indeed has a particular intention but

⁹⁸ Or perhaps a belief about what constitutes achieving E. I'll leave this complication to the side in what follows.

⁹⁹ It's worth mentioning that despite its means-ends form, this conception is neutral about what is involved in setting various ends. Some ends are surely reasoned to on the basis of higher-order ends, but even the highest-order ends might simply be caused or they might result from some other form of deliberation.

rather the very *intention* to do so. Perhaps this can constitute knowledge of what to do, but even so, this is a very special case of knowledge, and so it is only a very special instance of (10). Consider, then, the minor premise. It shouldn't be disputed in this context that we *can* have knowledge about the best means to an end. So another instance of (10) would be one's appropriate reliance on a means-end belief as the minor premise in a piece syllogistic practical reasoning. But this too is a special case of knowledge, and so it is only a relatively special instance of (10).

But what about much more ordinary and clearly knowable propositions, such as Williamson's example of the proposition that his door is locked? What our present reflections indicate is that, given the classical syllogistic conception of practical reasoning, it isn't at all clear how to make sense of relying on such a proposition in practical reasoning. For most ordinary (and clearly knowable) propositions are not the right form to be the major or minor premises in a practical syllogism. They set no ends. And they don't relate means to ends. This means that either all of these ordinary and clearly knowable propositions provide counterexamples to (10) or what it means to be "entitled to rely on her belief that *q* in her practical reasoning" must be something other than that its expression can serve as the major or minor premise in a practical syllogism.

So I think Williamson must be conceiving of practical reasoning more broadly than what gets formalized in the classical practical syllogism. The most natural suggestion is that he is thinking of practical reasoning as comprising not only the rational transition between an intention to achieve a particular end and a means-ends belief about how to achieve that end but also as the rational transitions which culminate in means-ends beliefs of the form P2. So, to take up his example of believing that the door is locked: this may figure as a premise in the following piece of reasoning:

(T1) The door is locked.

(T2) If the door is locked, then the quickest way to get to work while having my house secure is to get in my car now.

(TC) Thus the quickest way to get to work while having my house secure is to get in my car now.

And this can combine in the natural way with the major premise of a practical syllogism:

(P1) I shall get to work as quick as possible while having my house be secure.

To reach the the practical conclusion:

(PC) I shall get in my car now.

If this is on the right track, then one way that a belief can serve as an appropriate premise in practical reasoning—as Williamson must be conceiving of it—is by being an appropriate premise in a piece of theoretical reasoning leading to a belief about the best means to some end. However, this seems to mean that there is nothing special about the connection between knowledge and entitlement to rely on a belief in *practical* reasoning. For it seems as if any belief could serve as the major premise of a piece of *theoretical* reasoning leading to a conclusion about the best means to some end. So, although Williamson may be right that there is “an intuitive link between knowledge and practical reasoning” this appears to be because there is an intuitive link between knowledge and reasoning, *tout court*. This encourages a broader statement than (10)—viz.:

(11) S knows that-*q* iff S is entitled to rely on her belief that *q* in her reasoning.

And, as an account of knowledge attributions, (11) says that in attributing knowledge, we are attributing a specific sort of entitlement—it’s the epistemic entitlement to rely on a belief in one’s reasoning.

In pursuit of a general account of the practical point of knowledge claims, (11) encourages viewing *relying on a belief in one’s reasoning* as an action that is inferentially linked to

knowledge claims. The idea would be that a knowledge claim of the form (3) inferentially supports the subject S herself in relying in her reasoning on the beliefs that comprise her knowledge. I think this is another plausible proposal linking knowledge claims to action. However, unfortunately, it cannot be the complete story. For the Williamson-inspired proposal forges a link between knowledge claims and justified actions only in the special case of first-personal present-tense knowledge claims of the form,

(12) I know that p.

This leaves out cases of third-personal or past-tensed knowledge claims. For it isn't clear on this account what sort of actions these should be thought to justify.

So far we've seen three accounts of the practical point of knowledge claims. In each case, I've interpreted them as providing an explanation of what actions are inferentially linked to knowledge claims. Craig's account links knowledge claims to treating someone as a good informant about a piece of information. Rosenberg's account links knowledge claims to treating someone as having legitimately closed certain questions in an ongoing inquiry of which we are, at least implicitly, a part. And the account inspired by Williamson links knowledge claims to relying on beliefs in one's reasoning. In each case, the accounts were found to be lacking in their generality. That is, for however well the accounts worked for some species of knowledge claims, there seemed to be instances of knowledge claims that could not be handled by the accounts.

What I want to propose now is that we view each of the three accounts as three (perhaps not all-inclusive) components of a more general account. The epistemic actions isolated by Craig, Rosenberg, and Williamson should be seen as three species of a genus of epistemic action that I shall technically dub *trusting someone about a piece of information*.¹⁰⁰ One species of trusting someone about a piece of information is when we treat him as a good informant. This is the

¹⁰⁰ To be sure, often we use the word 'trust' to denote an attitude rather than an action. Indeed, for my purposes. All that is important is that a commitment to trusting someone has the inferential character of a practical rather than theoretical commitment.

insight of Craig's account. Another species of trusting someone about a piece of information is when we trust ourselves or a member of our epistemic community by taking the knower to have done everything necessary to close a question in an ongoing inquiry of which we are, at least implicitly, a part. This is the insight of Rosenberg's account. Both Craig's and Rosenberg's accounts had a hard time explaining the practical point of knowledge claims which attribute knowledge to someone who is neither a member of one's own community of inquiry nor necessary as an informant in the ongoing inquiry. Initially, the account inspired by Williamson doesn't appear to help with these cases either, since it appeared to link knowledge claims only to the subject's action of relying on a belief in her practical reasoning. However, I think we should recognize as another species of the action of trusting someone about a piece of information the action of treating someone as entitled to use particular premises in his reasoning. This will involve, *inter alia*, coming to believe what he believes, allowing others to defer justificatory responsibility to him, treating him as capable of meeting this responsibility, etc. These are all actions that I want to treat as components of the generic action of trusting someone about a piece of information.

If this line of thought is on the right track, then we have isolated a genus of actions—the actions which comprise trusting someone about a piece of information—that appear to be inferentially linked to knowledge claims. Thus, the inferentialist expressivist account of knowledge claims will say that these claims have this action in their downstream inferential potential in the special way definitive of expressions of practical commitments. That is, knowledge claims express_(c) practical commitments because they can serve as the sole practical premise in a piece of valid practical reasoning that culminates in an intention to perform the action of trusting someone about a piece of information. I think this represents a version of epistemic expressivism. Since it turns on a contrast between what is expressed by ordinary descriptive claims and what is expressed by knowledge claims. However, it is a novel form of expressivism in that it is not a version of noncognitivism. That is to say that the relevant

expressivist contrast is not drawn in terms of a contrast between cognitive and noncognitive psychological states putatively expressed by different sorts of claims. Instead, the relevant expressivist contrast is drawn in inferentialist terms—in terms, that is, of the types of inferences the claims can serve in.

Novelty is perhaps a virtue of a philosophical account, but it remains to be shown that this inferentialist expressivist account of knowledge claims can overcome the problems besetting the descriptivist and noncognitivist accounts considered up to this point. So I want to turn to this presently.

In the previous chapter, I identified three complex challenges facing various descriptivist accounts of the meaning of knowledge claims. I suggested there that we take meeting these challenges as adequacy conditions on an acceptable account of the meaning of knowledge claims. We saw earlier in this chapter that, *modulo* certain concerns, the sort of epistemic noncognitivism that I modeled on Gibbard's (1990) norm-expressivist account of ethical claims does a better job meeting these challenges than any of the descriptivist accounts explored in the previous chapter. Now, however, I want to argue that the inferentialist expressivist account of knowledge claims just sketched does even better than epistemic noncognitivism meeting the three challenges, while avoiding the concerns it raised.

The first challenge was Austin's challenge, which was, in effect, to avoid an account of the meaning of knowledge claims which makes ordinary epistemic discourse systematically and widely mistaken. This is made difficult by the fact that in some contexts we attribute knowledge quite freely, while in other contexts we are much more conservative about whether we'll say that someone has knowledge. To treat one or the other of these general features of ordinary epistemic discourse as incorrect seems to involve us in the posit of systematic and widespread error. Such mistakes are of course possible, but any account of the meaning of a term which

implicates such mistakes owes a very good explanation of them. Of the descriptivist accounts, contextualism seemed best equipped to meet this challenge. And the norm-expressivist account sketched above was able to piggyback on the contextualist strategy for meeting this challenge.

We worried, however, that this might be an ad hoc move. Sure—a critic may say—the norm-expressivist account meets Austin’s challenge as well as the contextualist, but it doesn’t really progress the explanation of why our knowledge attributions vary in an apparently context sensitive way. Why do we accept some epistemic norms in some situations and others in other situations?

I think the inferentialist expressivist account of knowledge claims advanced above has a way to incorporate the insight of contextualism vis-à-vis Austin’s challenge, while improving the explanation of why knowledge attributions apparently vary in the contextualist way. For, on this account, knowledge claims are construed as expressing a special sort of practical commitment precisely in that they can justify the epistemic action of trusting someone about a piece of information. As we’ve already seen, there are multiple species of this action, and, on many of them, it’s fairly easy to see practical considerations that would affect the conditions under which it makes sense to trust someone about a piece of information. Consider, for instance, the species of knowledge attributions we found to be handled well by Craig’s account of knowledge claims as flagging good informants. Whether it is reasonable to treat someone as a good informant will surely depend on the purposes for which we need the information. So, to take a now standard contextualist example, if we are looking for information about whether a flight has a stopover in Chicago, who counts as a good informant—and thus who counts as knowing whether the flight has a stopover in Chicago—will vary with the purpose of our seeking the information. If we’re just trying to figure out whether we’ll have to waste an hour in the Chicago airport, then someone who overheard the flight attendant say the flight has a stopover might suffice as a good informant. However, if we’re trying to plan a high-powered face-to-face meeting in the Chicago

airport between two executives, then we might require an informant with better epistemic credentials.

This is just an example, but the form of argument is relatively clear. Once we see knowledge claims as expressing practical commitments capable of inferentially supporting the action of trusting someone about a piece of information, whatever this sort of trusting comes to, variations in the ordinary practice of attributing knowledge will be explicable in terms of variations in the conditions under which it is reasonable to trust someone about a piece of information. In this way, the inferentialist expressivist account of knowledge captures the insight behind the contextualist response to Austin's challenge, while offering new resources for understanding why ordinary epistemic discourse exhibits the sorts of variations that it does.

The second challenge emerging out of our discussion in the previous chapter was the challenge of the paradox. The challenge was to provide a coherent response to the AFC-paradoxes. Of the descriptivist accounts, the contextualist account does the best job meeting this challenge because it offers a unified response to the first- and third-personal versions of these paradoxes, which doesn't have to explain away intuitive commitment to any of the elements of the paradoxes. The core idea is to grant that skeptical arguments are sound in the context in which they are presented but then insist that their conclusion is false when uttered in other contexts, which is why we think the argument must be unsound.

We worried that the contextualist response to the AFC-paradoxes might be an ad hoc maneuver because it could over-generalize to "solve" *any* paradox, and so it cannot explain why philosophers have been particularly troubled by the AFC-paradoxes. The noncognitivist version of epistemic expressivism sketched above offers essentially the same solution to the paradoxes, but it brought new resources to bear on worries about the solution being ad hoc. The idea was to use the thesis that knowledge claims express states of norm-acceptance to explain how the skeptical arguments are as attempts to change the epistemic norms that we accept. Moreover,

this explanation generalizes only to cases of language plausibly thought to express the acceptance of norms.

Although this provides an explanation of what is going on when we find the AFC-paradoxes paradoxical, as long as states of norm-acceptance are thought to be noncognitive states, the norm-expressivist proposal fails to explain how it is that we *reason* from the premises of the skeptical argument to its counterintuitive conclusion. Above I suggested that this renders epistemic norm-expressivism problematic. For the epistemic norm-expressivist owes an account of why these noncognitive attitudes of norm-acceptance vary in the way that they do. We saw that Gibbard (1990) attempts such an account in the ethical case, but no one has pursued a similar line in the epistemic case. Perhaps it remains to be done, but what is important for my purposes is that the inferentialist expressivist account of knowledge claims I've sketched can overcome this residual worry about epistemic norm-expressivism while providing essentially the same response to the challenge of the paradox.

The inferentialist expressivist can argue that the paradoxes are generated by shifts in our (epistemic) practical commitments across contexts. Skeptical considerations may indeed lead us to think that it is unreasonable to trust someone about a piece of information, even if we thought it perfectly reasonable to trust before entertaining the skeptical considerations. In such a case, the inferentialist expressivist proposal predicts that we will be less willing to attribute knowledge after entertaining the skeptical considerations.

Moreover, by treating knowledge claims as expressive of practical commitments that can justify the action of trusting someone about some piece of information, the inferentialist expressivist can also, where appropriate, rationalize contextual shifts in our willingness to attribute knowledge. It's not just a matter of what epistemic standards happen to be in play (as the contextualist would have it) nor of what epistemic norms the speaker happens to accept (as the norm-expressivist would have it), rather it is a matter of whether one takes the actions

comprising trusting someone about a piece of information to be justified. And it's easy to see why the answer to this question may vary depending on what hangs on so trusting.

The third challenge that emerged from our discussion in the previous chapter was the problem of unexplained semantic blindness. All of the descriptivist accounts posited some form of semantic blindness either to ordinary speakers or to philosophers. And in each case, it was far from clear how to explain the existence of such semantic blindness. The invariantist accounts were forced to posit semantic blindness at the level of first-order knowledge claims. Accounts that allow for contextual shifts the acceptability of an attribution of knowledge have a way to escape this posit of semantic blindness. But the contextualist account had then to posit semantic blindness about the dialectical connections between knowledge claims made in different contexts. The problem was that not only do ordinary speakers seem to attribute and deny knowledge (in different contexts) to the same subject of the same proposition based on the same evidence, they also seem prone to judge that such knowledge claims stand in some sort of conflict. Indeed, the intuition of conflict appears to be what funds the sense of paradoxicality in the AFC-paradoxes.

We saw above that the norm-expressivist about of knowledge claims has a way to meet this challenge. He follows the contextualist explanation of contextual shifts in first-order knowledge claims but he captures the meta-level intuition of conflict not at the level of descriptive content but at the level of the expression of the acceptance of norms. The idea was that a knowledge claim which expresses the acceptance of one norm N_1 can stand in conflict with a knowledge claim which expresses the acceptance of another norm N_2 , as long as one cannot simultaneously accept N_1 and N_2 . We might worry, however, that this sort of disagreement in noncognitive attitude is not exactly the sort of disagreement we need in order to capture the intuition of dialectical connection—at least not in the case of a “paradox”, which, after all, is supposed to include mutually inconsistent claims.

Like the norm-expressivist account of knowledge claims, the inferentialist expressivist account allows for contextual shifts in the acceptability of attributions of knowledge. So, at the level of first-order knowledge claims, it falls to Austin's challenge as the invariantist did.¹⁰¹ And at the level of second-order intuitions about conflicts between knowledge claims, I think it also captures the intuition of dialectical connection better than norm-expressivism. According to inferentialist expressivism, two knowledge claims will be in conflict if they express conflicting practical commitments. This is similar to the norm-expressivist strategy for capturing the dialectical connections between knowledge claims, but I think there are two elements of the inferentialist expressivist story which provide for an even more satisfactory account. First, it is clearer what sort of conflict is at issue between a knowledge claim and its negation. We already saw this sort of conflict when discussing ethical claims. One who says, "I'm going to bed," expresses a practical commitment to go to bed; and one who says, "I'm going to stop you," expresses a practical commitment that is in *practical* conflict with the commitment to go to bed. Even though this is not logical contradiction, it is a clear and intuitive sort of conflict. Moreover, it is clearly irrational to hold practical commitments that stand in this sort of conflict. Second, as in the ethical case, there seems to be some pressure to think of knowledge claims as expressing *collective* practical commitments. As we saw, we can model their content in terms of 'we-shall's. The idea was to treat,

(3) S knows that p,

as expressing a commitment with the content,

(3") We shall trust S with regards to p.

¹⁰¹ In this sense, it incorporates a form of contextualism about the truth-conditions of first-order knowledge claims, and so it will treat as semantically blind anyone who denies that the truth conditions of knowledge claims are implicitly relative. However, what is important in this debate is to get an account of the meaning of knowledge claims that can explain the cross context shifts in the attributability of knowledge. And the inferentialist expressivist can do this as well as the contextualist for first-order knowledge attributions.

It is not always obvious who should count as the ‘we’, but, in some cases, something like Rosenberg’s “community of inquiry” seems correct. So, in such cases, if one member of the community of inquiry says, e.g.,

(13) Holmes knows that the butler was at home during the murder,

this might be thought to express a commitment with the content,

(13") We shall trust Holmes with regards to the information that the butler was at home during the murder.

And, if another member of the community of inquiry says,

(14) Holmes does not know that the butler was at home during the murder,

this, in turn, might be thought to express a commitment with the content,

(14") We shall not trust Holmes with regards to the information that the butler was at home during the murder.

And clearly where the scope of the ‘we’s is the same we have a practical conflict between (13) and (14). Both practical commitments cannot be carried through. I think these considerations show that the inferentialist expressivist account of knowledge claims is even better equipped to handle the problem of unexplained semantic blindness than the norm-expressivist account, and since the norm-expressivist account did a better job than all of the descriptivist accounts, the inferentialist expressivist account comes out on top on this score.

All of this bodes well for the inferentialist expressivist account of knowledge claims, although a would-be epistemic norm-expressivist might hold out hope for a substantive positive account of why we express noncognitive attitudes of epistemic norm-acceptance in making knowledge claims. With such an account, I think it would be unclear which of norm-expressivism and inferentialist expressivism better meets the three challenges. However, our discussion of epistemic norm-expressivism generated a fourth challenge. This was the challenge of creeping minimalism and valid reasoning. The issue is complex because it arises deep in the

dialectic about noncognitivist accounts, but the basic problem is that it looks as if the noncognitivist can save the appearance of truth-evaluability and valid reasoning in his target area of discourse only by adopting a minimalist stance towards notions like truth and belief. That is, for example, the epistemic norm-expressivist must grant that knowledge claims are truth-evaluable and expressive of beliefs but, then, go on to argue that this is the case only in a minimalist sense of ‘truth-evaluable’ and ‘belief’. But the problem with this is that it then becomes unclear what the operative difference is for the noncognitivist between areas of discourse that are interpreted cognitively and those that are interpreted noncognitively. What, that is, distinguishes epistemic discourse from descriptive discourse so that the former expresses noncognitive states while the latter expresses cognitive states? There are some options for the epistemic noncognitivist, but as long as the distinction must be made in terms of ‘truth’ and its semantic relatives or in terms of ‘belief’ and its psychological relatives, each of the options seems to suffer the same problems as ‘truth’ and ‘belief’. This is creeping minimalism.

This is a familiar dialectic from recent metaethical debate, which I explored in detail in chapter 2. There I argued that one of the key advantages of inferentialist expressivism about ethical claims is that it is an expressivist account, which, even in the face of creeping minimalism, has the resources to mark out an intuitive and theoretically useful expressive contrast. The idea is to say that all of the relevant sentences $\text{express}_{(s)}$ propositions, and utterances of these sentences $\text{express}_{(p)}$ beliefs (plus, perhaps, other mental states), but one who makes the claim $\text{expresses}_{(c)}$ a commitment. And this commitment can be of one of two types—*theoretical* or *practical*. This marks out a clear difference between theoretical and practical claims. The former $\text{express}_{(c)}$ theoretical commitments, while the latter $\text{express}_{(c)}$ practical commitments.

In chapter 2, I explored in some detail ways we might spell out this distinction further and the sorts of semantic, epistemological, and ontological consequences it has. What is important

here is that an inferentialist expressivist about knowledge claims can avail himself of this distinction to meet the challenge of creeping minimalism and valid reasoning that besets epistemic norm-expressivism. For he can say that whatever we say at the level of ‘truth’ and ‘belief’, when it comes to theorizing about the commitments one expresses_(c) in making various claims, there is a clear difference between claims that express_(c) a theoretical commitment and claims that express_(c) a practical commitment. And the inferentialist expressivist view is that knowledge claims express a practical commitment in that they can serve as the sole practical premise in a piece of valid practical reasoning leading to the action of trusting someone about a piece of information. As in the ethical case, Frege-Geach inspired worries about whether an epistemic expressivist can recognize the embedding of knowledge claims in patterns of valid inference should vanish once we draw the relevant expressive contrast in explicitly inferentialist terms. We no longer conceive of the expressive contrast as one of whether the type of mental state is cognitive or noncognitive, as the epistemic-norm expressivist did; rather we conceive of the expressive contrast as one of whether the type of commitment expressed is theoretical or practical. Since both theoretical and practical commitments are defined in terms of their roles in inferences, it should be clear that claims which express them can equally figure in patterns of valid inference.

This concludes my presentation and defense of inferentialist expressivism. In metaethical debate, we’ve seen that the position offers a new option which meets the challenges threatening all descriptivist accounts; moreover, the position meets the key challenge threatening contemporary quasi-realist noncognitivist accounts. Often I am surprised that metaethicists fail to ask how their favored account of the meaning of ethical claims would play out when

transposed to an account of the meaning of epistemic claims. To be sure, the implausibility of a metaethical account transposed into the metaepistemological arena is not a fatal flaw—it may just highlight a key difference between ethical and epistemic discourse. However, since both areas of discourse are paradigms of normative discourse, and since recent metaethical and metaepistemological debate exhibits certain obvious parallels, I think an account that works in both arenas gains a sort of plausibility not possessed by other accounts. This encouraged the investigation of chapters 3-4. In the course of this investigation, we've seen that the inferentialist expressivist account extends well to metaepistemology and is able to meet the key challenges threatening all descriptivist accounts; moreover, as in the metaethical case, the position meets the key challenge threatening the most plausible version of epistemic norm-expressivism.

BIBLIOGRAPHY

- Alston, William. 1965. "Expressing," in *Philosophy in America*. Max Black ed. Ithaca: Cornell University Press.
- Alston, William. 1989. *Epistemic Justification*. Ithaca: Cornell University Press.
- Anscombe, G. E. M. 1957. *Intention*. Oxford: B. Blackwell.
- Armstrong, David. 1982. "Metaphysics and Supervenience." *Critica*, 14, pp. 3-18.
- Austin, J. L. 1953. "How to Talk—Some Simple Ways." *Proceedings of the Aristotlean Society*, 53, pp. 227-46.
- Austin, J. L. 1962. *Sense and Sensibilia*. Oxford: Oxford University Press.
- Ayer, A. J. 1936/1946. *Language, truth and logic*. London: V. Gollancz ltd.
- Baldwin, Tom. 2004. "George Edward Moore." *The Stanford Encyclopedia of Philosophy* (Summer 2004 Edition).
- Bar-On, Dorit. 2004. *Speaking my mind : expression and self-knowledge*. Oxford; New York: Clarendon Press ; Oxford University Press.
- Bar-On, Dorit, and Keith Simmons. forthcoming. "The use of force against deflationism: assertion and truth," in *The Illocutionary Role of the Concept of Truth*. D. Greimann and G. Siegart ed: de Gruyter.
- Blackburn, Simon. 1984. *Spreading the word : groundings in the philosophy of language*. Oxford [Oxfordshire]; New York: Clarendon Press ; Oxford University Press.
- Blackburn, Simon. 1988/1993. "Attitudes and Contents," in *Essays in Quasi-Realism*. New York: Oxford University Press.
- Blackburn, Simon. 1993. *Essays in quasi-realism*. New York: Oxford University Press.
- Blackburn, Simon. 1998. *Ruling passions : a theory of practical reasoning*. Oxford; New York: Clarendon Press ; Oxford University Press.
- Blackburn, Simon and Keith Simmons. 1999. *Truth*. Oxford: Oxford University Press.
- BonJour. 1978. "Can Empirical Knowledge Have a Foundation." *American Philosophical Quarterly*, 15, pp. 1-13.
- Boyd, Richard. 1988. "How to Be a Moral Realist," in *Essays on Moral Realism*. Geoffrey Sayre-McCord ed. Ithaca, NY: Cornell University Press.
- Brandom, Robert. 1994. *Making it explicit : reasoning, representing, and discursive commitment*. Cambridge, Mass.: Harvard University Press.

- Brandom, Robert. 2000. *Articulating reasons : an introduction to inferentialism*. Cambridge, Mass.: Harvard University Press.
- Bratman, Michael. 1987. *Intention, plans, and practical reason*. Cambridge, Mass ; London: Harvard University Press.
- Brink, David Owen. 1989. *Moral realism and the foundations of ethics*. Cambridge ; New York: Cambridge University Press.
- Capellan, Herman, and Ernest Lepore. 2003. "Context Shifting Arguments." *Philosophical Perspectives*, 17, pp. 25-50.
- Chrisman, Matthew. 2005a. "Review of Allan Gibbard's Thinking How to Live." *Ethics*, 115:2.
- Chrisman, Matthew. 2005b. "Review of Russ Shafer-Landau's Moral Realism." *Ethics*, 116:1, pp. 250-55.
- Chrisman, Matthew. forthcoming. "From Epistemic Contextualism to Epistemic Expressivism." *Philosophical Studies*.
- Cohen, Stewart. 1989. "How to Be a Fallibilist." *Philosophical Perspectives*, 2, pp. 91-123.
- Cohen, Stewart. 1999. "Contextualism, Skepticism, and the Structure of Reasons." *Nous*, 33: Supplement: Philosophical Perspectives, 13, Epistemology, pp. 57-89.
- Cohen, Stewart. 2005. "Knowledge, Speaker and Subject." *The Philosophical Quarterly*, 55:219, pp. 199-212.
- Copp, David. 2001. "Realist-Expressivism: A Neglected Option for Moral Realism." *Social Philosophy and Policy*, 18, pp. 1-43.
- Craig, Edward. 1990. *Knowledge and the state of nature : an essay in conceptual synthesis*. Oxford: Clarendon Press.
- D'Arms, Justin and Daniel Jacobson. 2000. "Sentiment and Value." *Ethics*, 110:4, pp. 722-48.
- Darwall, Stephen L., Allan Gibbard, and Peter Albert Railton. 1997. *Moral discourse and practice : some philosophical approaches*. New York: Oxford University Press.
- Darwall, Stephen L. 2003. "Moore, Normativity, and Intrinsic Value." *Ethics*, 113:3, pp. 468-90.
- Davidson, Donald. 1978/2001. "Intending," in *Essays on Actions and Events*. Oxford; New York: Oxford University Press.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research*, 52, pp. 913-29.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *The Philosophical Review*, 104:1, pp. 1-52. Reprinted in DeRose and Warfield (1999).
- DeRose, Keith and Ted A. Warfield. 1999. *Skepticism : a contemporary reader*. New York: Oxford University Press.

- DeRose, Keith. 2002. "Assertion, Knowledge, and Context." *Philosophical Review*, 111:2, pp. 167-203.
- DeRose, Keith. 2004. "Single Scoreboard Semantics." *Philosophical Studies*, 119:1 - 2, pp. 1-21.
- DeRose, Keith. 2005. "The Ordinary Language Basis for Contextualism and the New Invariantism." *The Philosophical Quarterly*, 55, pp. 172-98.
- DeRose, Keith. forthcoming. "'Bamboozled by Our Own Words': Semantic Blindness and Some Objections to Contextualism." *Philosophy and Phenomenological Research*.
- Divers, J. & Miller, A. 1994. "Why Expressivists About Value Should Not Love Minimalism About Truth." *Analysis*, 54, pp. 12-19.
- Dreier, James. 1990. "Internalism and Speaker Relativism." *Ethics*, 101:1, pp. 6-26.
- Dreier, James. 1996. "Expressivist Embeddings and Minimalist Truth." *Philosophical Studies*, 83:1, pp. 29-51.
- Dreier, James. 1999. "Transforming Expressivism." *Nous*, 33:4, pp. 558-72.
- Dreier, James. 1999. "Transforming Expressivism." *Nous*, 33:4, pp. 558-72.
- Dreier, James. 2004. "Meta-Ethics and the Problem of Creeping Minimalism." *Philosophical Perspectives*, 18, pp. 23-44.
- Dretske, Fred I. 1970. "Epistemic Operators." *The Journal of Philosophy*, 67:24, pp. 1007-23.
- Dretske, Fred I. 1981. *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Edwards, Paul. 1955. "Bertrand Russell's Doubts about Induction," in *Logic and Language*. A. Flew, ed. Oxford: Oxford University Press.
- Feldman, Richard. 1995. "In Defence of Closure." *The Philosophical Quarterly*, 45:181, pp. 487-94.
- Feldman, Richard. 1999. "Coherence, Contextualism, and Skepticism." *Feldman, Richard, Nous-Supplement: Philosophical Perspectives*, 13, *Epistemology*:33, pp. 91-114.
- Field, Harty. 1987. "The Deflationary Conception of Truth," in *Fact, Science, and Morality; Essays on A. J. Ayer's Language, Truth, and Logic*. C. Wright Graham Macdonald, and A. J. Ayer ed. New York, Ny: Blackwell.
- Field, Hartry. 1998. "Epistemological Nonfactualism and the a Prioricity of Logic." *Philosophical Studies*, 92:1 - 2, pp. 1-24.
- Field, Hartry. 2001. "Disquotational Truth and Factually Defective Discourse," in *Truth and the Absence of Fact*. New York, NY: Oxford University Press.
- Fine, Kit. 2001. "The Question of Realism." *Philosophers' Imprint*, 1:1, pp. 1-30.
- Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University

Press.

- Frankena, William. 1968. "Obligation and Value in the Ethics of G. E. Moore," in *The Philosophy of G. E. Moore*. Paul Aurther Schlipp ed. La Salle, Ill: Open Court.
- Frege, Gottlob. 1879. *Begriffsschrift, eine der arithmetischen Nachgebildete Formelsprache des reinen Denkens*. Halle a/S.: L. Nebert.
- Frege, Gottlob and Michael Beaney. 1997. *The Frege reader*. Oxford ; Cambridge, Mass.: Blackwell Publishers.
- Geach, P. T. 1965. "Assertion." *The Philosophical Review*, 74:4, pp. 449-65.
- Gibbard, Allan. 1985. "Moral Judgment and the Acceptance of Norms." *Ethics*, 96:1, pp. 5-21.
- Gibbard, Allan. 1986. "An Expressivistic Theory of Normative Discourse." *Ethics*, 96:3, pp. 472-85.
- Gibbard, Allan. 1990. *Wise choices, apt feelings : a theory of normative judgment*. Oxford: Clarendon Press.
- Gibbard, Allan. 1992. "Moral Concepts: Substance and Sentiment." *Philosophical Perspectives*, 6:Ethics, pp. 199-221.
- Gibbard, Allan. 2003. *Thinking how to live*. Cambridge, Mass.: Harvard University Press.
- Grice, H. P. 1989. *Studies in the way of words*. Cambridge, Mass.: Harvard University Press.
- Gupta, Anil and Nuel D. Belnap. 1993. *The revision theory of truth*. Cambridge, Mass.: MIT Press.
- Haidt, Jonathan. 2001. "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Review*, 108, pp. 814-34.
- Hale, Bob. 1993. "Can There Be a Logic of Attitudes," in *Reality, Representation, and Projection*. John Haldane, and Crispin Wright ed. Oxford: Oxford University Press.
- Hare, R. M. 1952. *The language of morals*. Oxford: Clarendon Press.
- Harman, Gilbert. 1975. "Moral Relativism Defended." *The Philosophical Review*, 84:1, pp. 3-22.
- Harman, Gilbert. 1977. *The nature of morality : an introduction to ethics*. New York: Oxford University Press.
- Harman, Gilbert and Judith Jarvis Thomson. 1996. *Moral relativism and moral objectivity*. Malden: Blackwell.
- Hofweber, Thomas. 1999. "Contextualism and the Meaning-Intention Problem," in *Cognition, Agency and Rationality*. Ernest Sosa Kepa Korta, and Xabier Arrazola ed: Kluwer Academic Publisher.
- Horwich, Paul. 1990. *Truth*. Oxford ; Cambridge, Mass.: B. Blackwell.

- Horwich, Paul. 1993. "Gibbard's Theory of Norms." *Philosophy and Public Affairs*, 22, pp. 67-78.
- Humberstone, Loyd. 1992. "Direction of Fit." *Mind*, 101, pp. 59-83.
- Hume, David and Tom L. Beauchamp. 1998. *An enquiry concerning the principles of morals : a critical edition*. Oxford; New York: Clarendon Press ; Oxford University Press.
- Hume, David and Tom L. Beauchamp. 1999. *An enquiry concerning human understanding*. Oxford ; New York: Oxford University Press.
- Hume, David, David Fate Norton, and Mary J. Norton. 2000. *A treatise of human nature*. Oxford ; New York: Oxford University Press.
- Jackson, Frank, Graham Oppy, and Michael Smith. 1994. "Minimalism and Truth Aptness." *Mind*, 103:411, pp. 287-302.
- Kant, Immanuel, Thomas E. Hill, and Arnulf Zweig. 2002. *Groundwork for the Metaphysics of Morals, Oxford Philosophical Texts*. Oxford ; New York: Oxford University Press.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research*, 66: 3.
- Klein, Peter D. 1981. *Certainty : a refutation of scepticism*. Brighton: Harvester.
- Korsgaard, Christine M. 1986. "Skepticism about Practical Reason." *The Journal of Philosophy*, 83:1, pp. 5-25.
- Korsgaard, Christine M. and Onora O'Neill. 1996. *The sources of normativity*. New York: Cambridge University Press.
- Kremer, Michael. 1987. "If is Unambiguous." *Nous*, 21:2, pp. 199-217.
- Kripke, Saul A. 1980. *Naming and necessity*. Cambridge, Mass.: Harvard University Press.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* (Australasian Journal of Philosophy), 74, pp. 549-67.
- Lycan, William G. 1988. *Judgement and justification*. Cambridge [Cambridgeshire] ; New York: Cambridge University Press.
- Lycan, William G. 2000. *Philosophy of language : a contemporary introduction*. London ; New York: Routledge.
- Macdonald, Graham, A. J. Ayer, and Crispin Wright. 1987. *Fact, science, and morality : essays on A.J. Ayer's Language, truth, and logic*. Oxford, UK ; New York, NY, USA: Blackwell.
- MacFarlane, John. 2005. "The Assessment Sensitivity of Knowledge Attributions," in *Oxford Studies in Epistemology*. Tamar Szabó Gendler, and John Hawthorne ed. Oxford: Oxford University Press, pp. 197-233.
- Mackie, J. L. 1977. *Ethics : inventing right and wrong*. Harmondsworth ; New York: Penguin.

- Majors, Brad. 2005. "Moral Discourse and Descriptive Properties." *The Philosophical Quarterly*, 55:220, pp. 475-94.
- McDowell, John. 1978. "Are Moral Requirements Hypothetical Imperatives?" *Proceedings of the Aristotlean Society*, supplementary volume.
- Mill, John Stuart. 1863. *Utilitarianism*. London: Parker, Son, and Bourn.
- Misak, Cheryl. 1991. *Truth and the end of enquiry*. Oxford: Clarendon Press.
- Moore, G. E. 1903. *Principia ethica*. Cambridge: University press.
- Moore, G. E. and Thomas Baldwin. 1993. *Principia ethica*. Cambridge ; New York, NY, USA: Cambridge University Press.
- Moran, Richard. 2001. *Authority and estrangement : an essay on self-knowledge*. Princeton, N.J.: Princeton University Press.
- Nagel, Thomas. 1970. *The possibility of altruism*. Oxford: Clarendon Press.
- Neta, Ram. 2003. "Skepticism, Contextualism, and Semantic Self-Knowledge." *Philosophy and Phenomenological Research*, 67, pp. 396 – 411.
- Neta, Ram. 2004. "Perceptual Evidence and the New Dogmatism." *Philosophical Studies*, 119:1 - 2, pp. 199-214.
- Nozick, Robert. 1981. *Philosophical explanations*. Cambridge, Mass.: Harvard University Press.
- O'Leary-Hawthorne, John, and Price, Huw. 1996. "How to stand up for non-cognitivists." *Australasian Journal of Philosophy*, 74: 275-292.
- Platts, Mark de Bretton. 1979. *Ways of meaning : an introduction to a philosophy of language*. London: Routledge and Kegan Paul.
- Platts, Mark. 1986. "Desire and Action." *Nous*, 20:2, pp. 143-55.
- Preyer, Gerhard and Georg Peter. 2005. *Contextualism in philosophy : knowledge, meaning, and truth*. Oxford ; New York: Clarendon ; Oxford University Press.
- Prinz, Jesse. forthcoming. *The emotional construction of morals*. New York: Oxford University Press.
- Pryor, James. 2001. "The Skeptic and the Dogmatist." *Noûs*, 34, pp. 517-49.
- Putnam, Hilary. 1960. "Minds and Machines," in *Dimensions of Mind*. Sidney Hook ed. New York: New York University Press, pp. 148-80.
- Putnam, Hilary. 1963. "Brains and Behavior," in *Analytical Philosophy*. R. J. Butler ed. Oxford: Basil Blackwell, pp. 211-35.
- Putnam, Hilary. 1975. *Mind, language, and reality*. Cambridge [Eng.] ; New York: Cambridge University Press.

- Railton, Peter. 1986. "Moral Realism." *Philosophical Review*, 95: 2.
- Regan, Donald H. 2003. "How to Be a Moorean." *Ethics*, 113:3, pp. 651-72.
- Richard, Mark. 2004. "Contextualism and Relativism." *Philosophical Studies*, 119:1 - 2, pp. 215-42.
- Ridge, Michael. 2003. "Moral Non-Naturalism." *The Stanford Encyclopedia of Philosophy (Spring 2003 Edition)*.
- Ridge, Michael. Forthcoming-a. "Ecumenical Expressivism: Finessing Frege." *Ethics*.
- Ridge, Michael. Forthcoming-b. "Sincerity and Expressivism." *Philosophical Studies*.
- Rosenberg, Jay F. 1974. *Linguistic representation*. Dordrecht, Holland ; Boston: D. Reidel Pub. Co.
- Rosenberg, Jay F. 2002. *Thinking about knowing*. Oxford: Clarendon Press.
- Ross, W. D. 1939. *Foundations of ethics; the Gifford lectures delivered in the University of Aberdeen, 1935-6*. Oxford: The Clarendon Press.
- Rysiew, Patrick. 2001. "The Context-Sensitivity of Knowledge Attributions." *Nous*, 35:4, pp. 477-514.
- Sartwell, Crispin. 1992. "Why Knowledge is Merely True Belief." *The Journal of Philosophy*, 89, pp. 167-80.
- Sayre-McCord, Geoffrey. 1988. *Essays on moral realism*. Ithaca, N.Y.: Cornell University Press.
- Sayre-McCord, Geoffrey. 1997. "'Good' On Twin Earth." *Philosophical Issues: Truth*, 8: Truth, pp. 267-92.
- Scanlon, Thomas. 1998. *What we owe to each other*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Schaffer, Jonathan. 2004. "From Contextualism to Contrastivism." *Philosophical Studies*, 119:1 - 2, pp. 73-103.
- Schiffer, Stephen. 1996. "Contextualist Solutions to Skepticism." *Proceedings of the Aristotelian Society*, 96, pp. 317-33.
- Schiffer, Stephen. 2004. "Skepticism and the Vagaries of Justified Belief." *Philosophical Studies*, 119:1 - 2, pp. 161-84.
- Searle, John R. 1962. "Meaning and Speech Acts." *The Philosophical Review*, 71:4, pp. 423-32.
- Sellars, Wilfrid. 1948/1980. "Inference and Meaning," in *Pure Pragmatics and Possible Worlds*. Sicha ed. Ridgeview: Ridgeview Publishing Co.
- Sellars, Wilfrid. 1953. "Inference and Meaning." *Mind*, 62:247, pp. 313-38.

- Sellars, Wilfrid. 1956. *Empiricism and the philosophy of mind*. Indianapolis, Ind.: Bobbs-Merrill.
- Sellars, Wilfrid. 1968. *Science and metaphysics: variations on Kantian themes*. London, New York: Routledge & K. Paul.
- Sellars, Wilfrid. 1969. "Language as Thought and as Communication." *Philosophy and Phenomenological Research*, 29:4, pp. 506-27.
- Sellars, Wilfrid and J. Sicha. 1980. *Pure pragmatics and possible worlds : the early essays of Wilfrid Sellars*. Reseda, Calif.: Ridgeview Pub. Co.
- Shafer-Landau, Russ. 2003. *Moral realism : a defence*. Oxford: Clarendon.
- Slovic, Paul, Fischhoff, Baruch, and Lichtenstein, Sarah. 1982. "Fact versus Fears: Understanding Perceived Risk," in *Choices, Values and Frames*. Daniel Kahneman, and Tversky, Amos ed. Cambridge: Cambridge University Press.
- Smart, J. J. C. 1959. "Sensations and Brain Processes." *The Philosophical Review*, 68:2, pp. 141-56.
- Smith, Michael. 1987. "The Humean Theory of Motivation." *Mind*, 96:381, pp. 36-61.
- Smith, Michael. 1994. "Why Expressivists About Value Should Love Minimalism About Truth." *Analysis*, 54, pp. 1-12.
- Smith, Michael. 1994. "Internal Reasons." *Philosophy and Phenomenological Research*, 55:1, pp. 109-31.
- Smith, Michael. 1995. *The moral problem*. Oxford: Blackwell.
- Smith, Michael. 2002. "Evaluation, Uncertainty, and Motivation." *Ethical Theory and Moral Practice*, 5, pp. 305-20.
- Smith, Michael. 2002. "Which Passions Rule?" *Philosophy and Phenomenological Research*, 65, pp. 157-63.
- Sosa, Ernest. 2004. "Relevant Alternatives, Contextualism Included." *Philosophical Studies*, 119:1 - 2, pp. 35-65.
- Stanley, Jason. 2004. "On the Linguistic Basis for Contextualism." *Philosophical Studies*, 119:1 - 2, pp. 119-46.
- Stanley, Jason. 2005. *Knowledge and practical interests*. Oxford ; New York: Clarendon Press ; Oxford University Press.
- Stevenson, Charles L. 1937. "Emotive Meaning of Moral Terms." *Mind*, 46, pp. 14-31.
- Stevenson, Charles L. 1944. *Ethics and language*. New Haven: Yale University Press.
- Stine, Gail. 1976/1999. "Skepticism, Relevant Alternatives, and Deductive Closure," in *Skepticism : a contemporary reader*. DeRose and Warfield ed. New York: Oxford

University Press.

- Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies*, 127:1, pp. 109-66.
- Stroud, Barry. 1984. *The significance of philosophical scepticism*. Oxford [Oxfordshire]: Clarendon Press.
- Sturgeon, Nicholas. 1985. "Moral Explanations," in *Morality, Reason, and Truth*. Copp and Zimmerman ed. Totowa: Rowman & Littlefield.
- Sturgeon, Nicholas L. 1992. "Nonmoral Explanations." *Philosophical Perspectives*, 6:Ethics, pp. 97-117.
- Sutton, Jonathan. 2005. "Stick to What You Know." *Nous* 39: pp. 359-396.
- Timmons, Mark. 1998. *Morality without foundations : a defense of ethical contextualism*. New York: Oxford University Press.
- Unger, Peter K. 1975. *Ignorance : a case for scepticism*. Oxford [Eng.]: Clarendon Press.
- Unger, Peter K. 1984. *Philosophical relativity*. Minneapolis: University of Minnesota Press.
- Velleman, J. David. 1992. "The Guise of the Good." *Nous*, 26:1, pp. 3-26. Reprinted in Velleman (2000).
- Velleman, J. David. 1996. "The Possibility of Practical Reason." *Ethics*, 106:4, pp. 694-726. Reprinted in Velleman (2000).
- Velleman, J. David. 1997. "How To Share An Intention." *Philosophy and Phenomenological Research: A Quarterly Journal*, 57:1, pp. 29-50. Reprinted in Velleman (2000).
- Velleman, James David. 2000. *The possibility of practical reason*. Oxford, New York: Clarendon ; Oxford University Press.
- Westermarck, Edward. 1932. *Ethical relativity*. London: Kegan Paul Trench Trèubner.
- Wiggins, David. 1991. *Needs, values, truth : essays in the philosophy of value*. Oxford: Blackwell.
- Williams, Bernard Arthur Owen. 1981. "Internal and External Reasons," in *Moral luck : philosophical papers, 1973-1980*. Cambridge: Cambridge University Press, pp. xiii, 173.
- Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford [England] ; New York: Oxford University Press.
- Williamson, Timothy. 2005a. "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge." *The Philosophical Quarterly*, 55:219, pp. 213-35.
- Williamson, Timothy. 2005b. "Knowledge, Context, and the Agent's Point of View," in *Contextualism in Philosophy*. Gerhard Preyer, and Georg Peter ed. New York; Oxford: Oxford University Press.

- Wolf, Susan. 1982. "Moral Saints." *The Journal of Philosophy*, 79:8, pp. 419-39.
- Wong, David B. 1984. *Moral relativity*. Berkeley: University of California Press.
- Wright, Crispin. 1992. *Truth and objectivity*. Cambridge, Mass. ; London: Harvard University Press.
- Yourgrau, Palle. 1983. "Knowledge and Relevant Alternatives." *Synthese*, 55, pp. 175-90.
- Zangwill, Nick. 1990. "Quasi-Quasi-Realism." *Philosophy and Phenomenological Research*, 50:3, pp. 583-94.